Gateway 840 VT-100 User's Guide

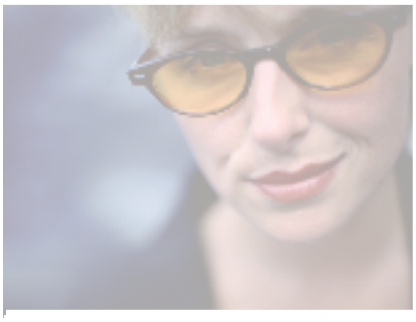

Installing | Getting Help | Gateway™

Contents

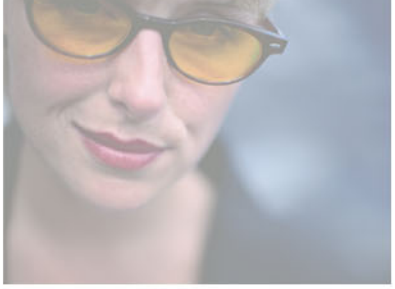

1 Introduction

This guide provides the operational and reference information to configure and manage the RAID controllers installed in your Gateway 840 Series storage system. The controllers are accessed through a VT-100 terminal using a menu-based user interface, connected directly to the RAID controllers.

This guide contains step-by-step procedures to access the VT-100 terminal utility, as well as perform configuration functions and enclosure monitoring.

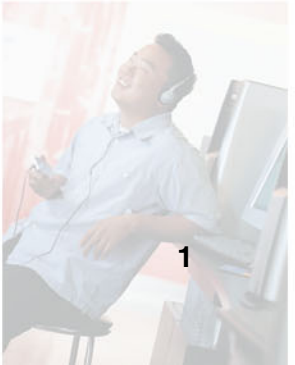

Features

Your Gateway 840 has a menu-based software utility designed to configure and manage the RAID controller in your new storage enclosure. It provides centralized management to control primary storage assets vital to ensuring maximum data reliability, network up-time, and system serviceability. It lets you manage the storage system by accessing the configuration utility locally.

This software guide provides the operational and reference information to configure and manage the SATA RAID controller installed in your Gateway 840, using this menu-based user interface.

The following are some of the major features of local RAID configuration utility.

- Simple, straight-forward user interface
- Menu-based interface that works with a wide range of operating systems
- Easy access and navigation
- Host LUN and LUN mapping configuration
- A comprehensive, non-volatile event log
- Useful offline or online diagnostics

Troubleshooting assistance

To help troubleshoot problems with your system, the Event Logs chapter, as well as on-line help, provides a list of all the events that can occur, along with a suggested cause and action to take. When a warning or error event occurs, you can review the event log, locate the suspect event, and see "Event Logs" on page 173 or help section to determine the possible cause or causes and the suggested action to take. Also see "Diagnostics" on page 199.

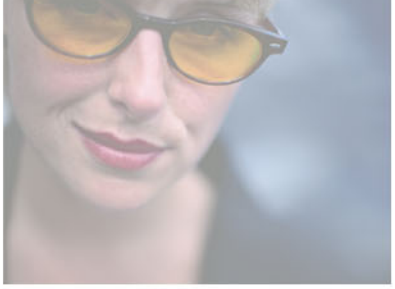

2 Accessing the Configuration Utility

This chapter provides information on accessing the configuration utility. Read this chapter to learn how to:

- Access the utility
- Use the VT-100 terminal
- Navigate the configuration utility
- View controller information

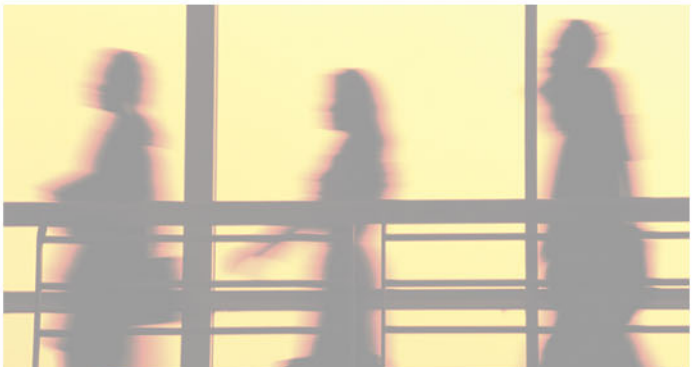

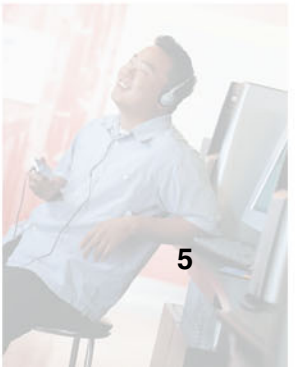

Accessing the utility

You can access the SATA RAID controller by using a VT-100 terminal connected to the RS-232 serial ports on the controller's rear panel. Configure your terminal using the procedures described in "Using the VT-100 terminal" on page 7.

Important

You can also choose to use StorView Storage Management software to manage and monitor the enclosure and SATA RAID Controller. Refer to the *Gateway 840 SATA RAID StorView User's Guide* for more information.

Using the VT-100 terminal

Access to the RAID controller is accomplished by using its onboard configuration utility with a VT-100 terminal or terminal emulation connected to the RS-232 serial port on the controller.

A SATA communication RS-232 cable (provided with the hardware) is used to connect the terminal to the serial port on the SATA RAID controller.

Configure your host system or terminal RS-232 port to use the following settings:

Setting	Value
Terminal Emulation	ANSI
Function	Terminal Keys
Fonts	Terminal
Translations	None
Columns	80

Set the communications parameters for the terminal program as follows:

Setting	Value
Baud Rate	115,200
Data Bits	8
Stop Bits	1
Parity	None
Flow Control	None
Connector	COM1 (typically)

To access the RAID controllers:

1. Connect the host terminal to the serial port on the controller using a null-modem serial cable (use the configuration information in the preceding tables).
2. Turn on the terminal and the Gateway 840 Series storage system.

3 From the computer or terminal connected to the RAID controller, start your terminal emulation software.

Important

Make sure that your terminal emulation software is set to use the correct COM port on your computer.

4 From the terminal window, press CTRL + W. The **Main Menu** opens.

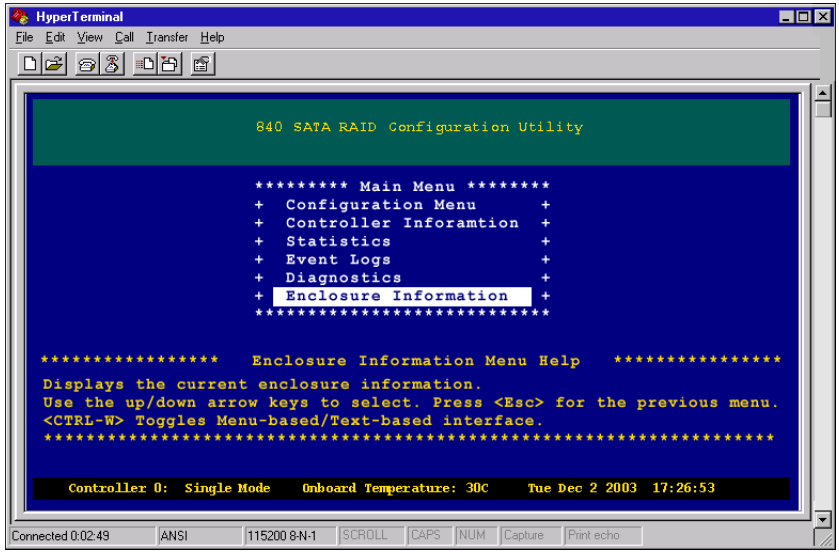

You can now perform all of the functions described in the following chapters. All steps begin from the Main Menu.

If an event or error occurs, you can review the log for the problem by examining the event log (see "Event Logs" on page 173).

Navigating the SATA RAID configuration utility

The SATA RAID configuration utility menu system can be navigated using the keys on your keyboard. The table below describes the primary navigation and selection methods.

To	Do this
Toggle interface from text-base to menu-based.	Press CTRL + W keys.
Return to the previous menu or screen (Exit).	Press the ESC key.
Move the selection.	Press the arrow keys.
Select a menu option.	Press the ENTER key.

Menu system

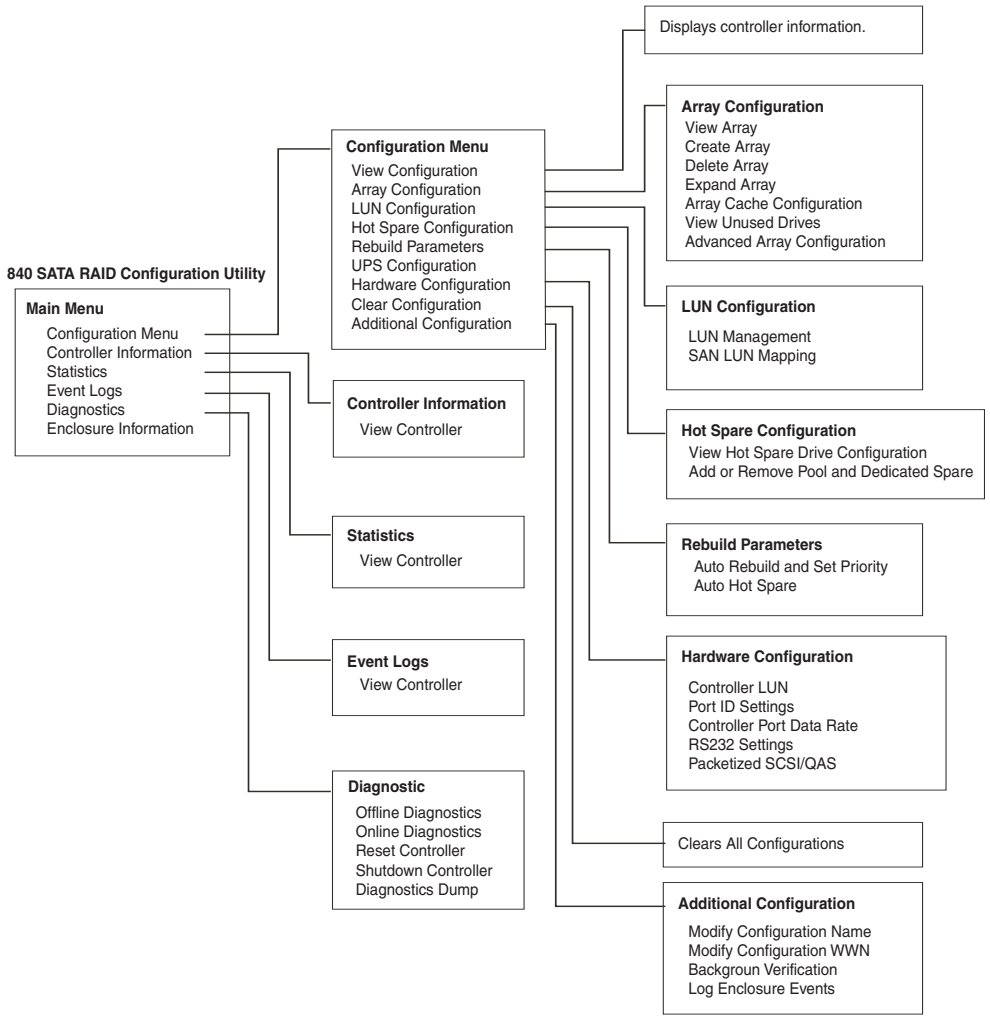
840 SATA RAID Configuration Utility
Main Menu
Configuration Menu
Controller Information
Statistics
Event Logs
Diagnostics
Enclosure Information
Configuration Menu
View Configuration
Array Configuration
LUN Configuration
Hot Spare Configuration
Rebuild Parameters
UPS Configuration
Hardware Configuration
Clear Configuration
Additional Configuration
Displays controller information.
Array Configuration
View Array
Create Array
Delete Array
Expand Array
Array Cache Configuration
View Unused Drives
Advanced Array Configuration
LUN Configuration
LUN Management
SAN LUN Mapping
Hot Spare Configuration
View Hot Spare Drive Configuration
Add or Remove Pool and Dedicated Spare
Rebuild Parameters
Auto Rebuild and Set Priority
Auto Hot Spare
Hardware Configuration
Controller LUN
Port ID Settings
Controller Port Data Rate
RS232 Settings
Packetized SCSI/QAS
Clears All Configurations
Additional Configuration
Modify Configuration Name
Modify Configuration WWN
Backgroun Verification
Log Enclosure Events
Controller Information
View Controller
Statistics
View Controller
Event Logs
View Controller
Diagnostic
Offline Diagnostics
Online Diagnostics
Reset Controller
Shutdown Controller
Diagnostics Dump

Menu system (continued)

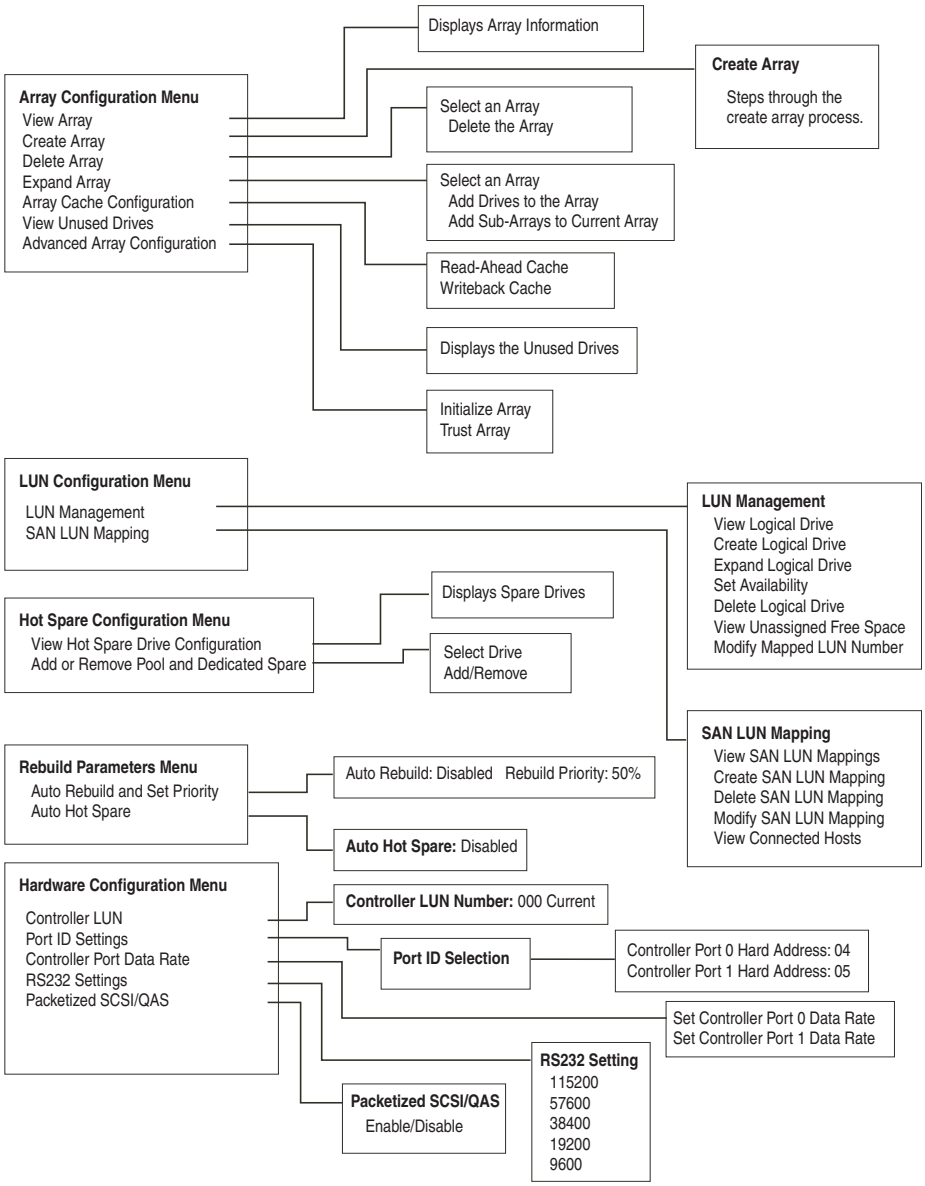

Menu system (continued)

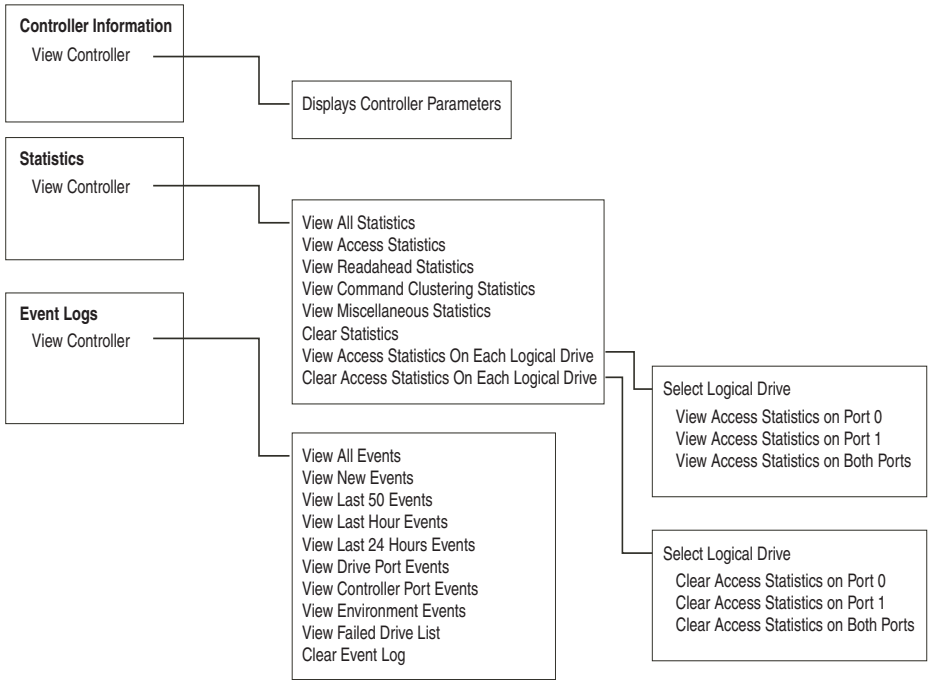

Viewing controller information

Detailed information on the RAID controller is available through the Main Menu, and includes:

- Controller World Wide Name (WWN)
- Firmware information
- Run-time information
- Temperature and voltage information

To view controller information:

1. After the controller(s) have completed booting, press **CTRL** + **E**, then press **CTRL** + **W** to switch to the menu-based mode. The *Main Menu* opens.

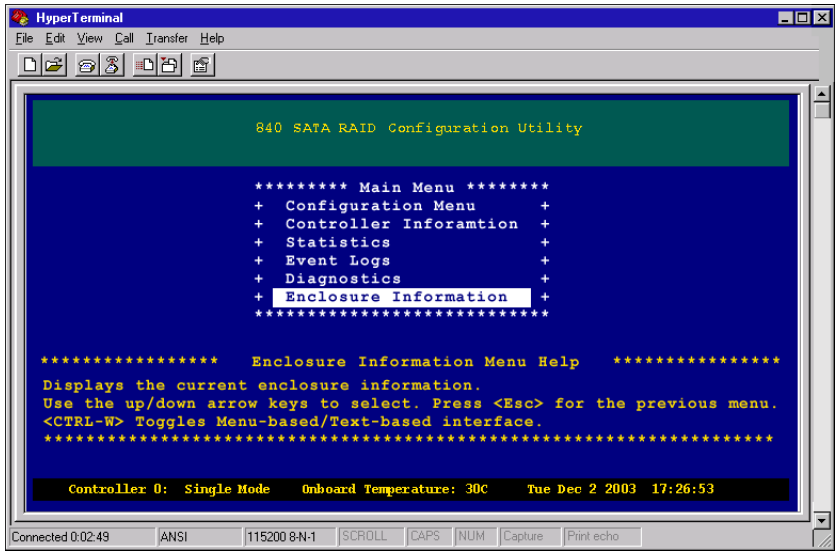

2 From the Main Menu select **Controller Information,** then press **ENTER**. The *Controller Information* screen opens.

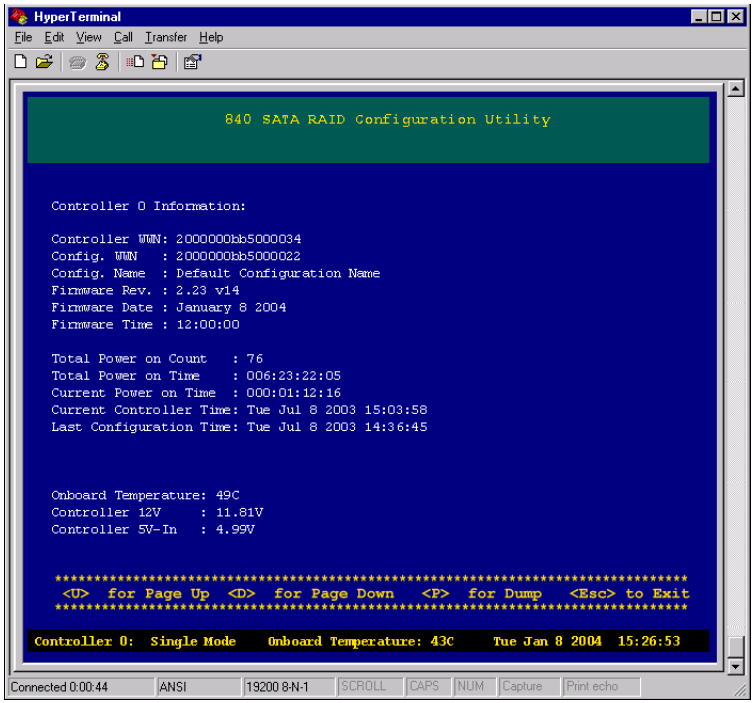

3 Select the controller you want to examine by using the spacebar to select the controller, then press **ENTER**. If only one controller is present, only one selection is listed.

The following is a brief summary of the information displayed.

Item	Description
Controller WWN	An 8-byte unique World Wide Name the controller uses for identification. This is unique to each controller, and preprogrammed.
Configuration WWN	An 8-byte World Wide Name the controller uses for identification to hosts. This can be identical to the Controller WWN, or can be different.
Config Name	A 64-byte ASCII name used to identify a specific controller configuration.

Item	Description
Firmware Revision	The firmware version in use by the controller.
Firmware Date/Time	Date and time of when the firmware was created.
Total Power on Count	The number of times this controller has been power cycled.
Total Power on Time	The total length of time this controller has been powered on.
Current Power on Time	The length of time since this controller has last been powered on or reset.
Current Controller Time	This is the current time on the controller.
Last Configuration Time	This is the time the controller was last configured.
Onboard Temperature	This is the current temperature as measured by the controller. If this value exceeds predefined limits, an event will be written into the event log.
Controller Voltage	Current readings for various controller and coprocessor (if applicable) voltages.

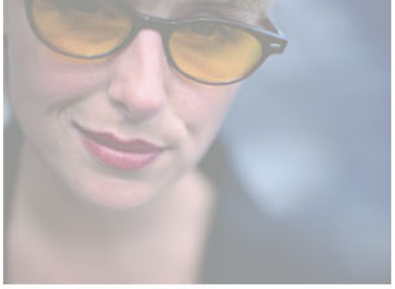

Creating Disk Arrays

3

This chapter provides information on creating disk arrays. Read this chapter to learn how to:

- Understand arrays
- Create arrays
- Configure array read-ahead and writeback cache
- Assign hot spare drives
- Delete an array
- Expand an array
- Trust an array

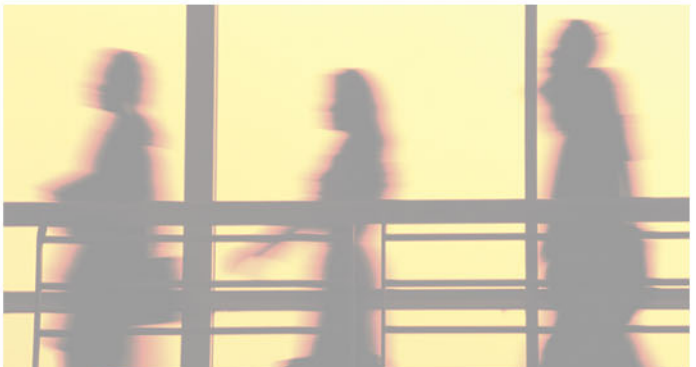

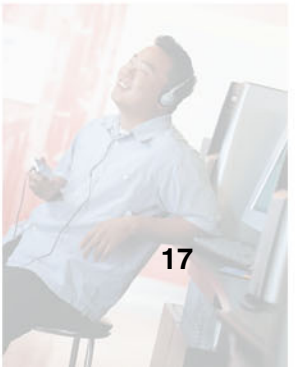

Overview

Configuring a RAID system requires some planning to make sure that you define the correct RAID levels and array options. It may be helpful to refer to the Topology chapter in the *Gateway 840 SATA RAID Enclosure User's Guide*.

This chapter will step you through the process to create the disk arrays. You will also find sections on deleting arrays, expanding arrays, and assigning hot spare drives. You are then directed to the next chapter for the procedures to define the logical drives, which makes the drive array(s) available to the operating system.

This manual assumes you have a basic understanding of RAID concepts and terminologies.

Understanding arrays

You can create an array at any time. The table below describes the drive requirements for each RAID level.

RAID Level	Minimum No. of Drives	Maximum No. of Drives
0	1	12
1	2	12
5	3	12
50	6	12
10 (Mirrored)	4	12

Important

Before you create more than one array, you must be sure that your host operating system supports multiple Logical Unit Numbers (LUNs). Most operating systems do support multiple LUNs, or have an option for it. If your operating system does not support multiple LUNs, the host will only be able to see one array at the first disk LUN.

Terminology

The following describes the terminology used when creating and managing arrays.

Term	Description
Array	A group of drives that are combined together to create a single large storage area. Up to 64 arrays are supported, each containing up to 12 drives per array. There is no limit for the drive size in the arrays.
Chunk Size	This is the amount of data that is written on a single drive before the controller moves to the next drive in the stripe.
Stripe Size	This is the number of data drives multiplied by the chunk size.
Cache Flush Array	This is the array that is used to automatically flush cache data in the situation where power has failed to some of the drives.

Initialization	RAID 5/50 arrays must have consistent parity before they can be used to protect data. Initialization writes a known pattern to all drives in the array. If you choose not to initialize an array, this is known as a “trusted array” and any drive failure will result in data corruption. It is possible to later perform a parity rewrite, which recalculates the parity based on the current data, thus ensuring the data and parity are consistent.
Reserved Capacity	In order to allow drives from a different family or manufacturer to be used as a replacement for a drive in an array, we recommend that a small percentage of the drive’s capacity be reserved when creating the array. This is user selectable, from 0 to 10 percent.
RAID Level 0	RAID 0 is defined as disk striping where data is striped or spread across one or more drives in parallel. RAID 0 is ideal for environments in which performance (read and write) is more important than fault tolerance or you need the maximum amount of available drive capacity in one volume. Drive parallelism increases throughput because all drives in the stripe set work together on every I/O operation. For greatest efficiency, all drives in the stripe set must be the same capacity. Because all drives are used in every operation, RAID 0 allows for single-threaded I/O only (i.e., one I/O operation at a time). Environments with many small simultaneous transactions (e.g., order entry systems) will not get the best possible throughput.
RAID Level 1	RAID 1 is defined as disk mirroring where one drive is an exact copy of the other. RAID 1 is useful for building a fault-tolerant system or data volume, providing excellent availability without sacrificing performance.
RAID Level 5	RAID 5 is defined as disk striping with parity where the parity data is distributed across with parity all drives in the volume. Normal data and parity data are written to drives in the stripe set in a round-robin algorithm. RAID 5 is multi threaded for both reads and writes because both normal data and parity data are distributed round-robin. This is one reason why RAID 5 offers better overall performance in server applications. Random I/O benefits more from RAID 5 than does sequential I/O, and writes take a performance hit because of the parity calculations. RAID 5 is ideal for database applications.

RAID Level 10	RAID 10 is defined as mirrored stripe sets (also known as RAID 0+1). You can build RAID 10 either directly through the RAID controller (depending on the controller) or by combining software mirroring and controller striping (called RAID 01).
RAID Level 50	This RAID level is a combination of RAID level 5 and RAID level 0. Individual smaller RAID 5 arrays are striped, to give a single RAID 50 array. This can increase the performance by allowing the controller to more efficiently cluster commands together. Fault tolerance is also increased, as one drive can fail in each individual array.
Sub-array	In RAID 50 applications, this is the name given to the individual RAID 5 arrays that are striped together. Each sub-array has one parity drive.

Optimization and drive selection for RAID 5 arrays

Typical RAID 5 implementations require a number of steps to write the data to the drives. In order to optimize your system performance based on the type of writes you expect in your operation, we have provided detailed information on optimizing the performance using full strip write operations in an appendix (see "Optimizing RAID 5 Write Performance" on page 227).

If you intend to setup a RAID 5 array and want to consider optimum performance, you will need to consider the number of drives, parity drives, and chunk size. You should review the information provided in "Optimizing RAID 5 Write Performance" on page 227. Additional information is provided at the appropriate step during configuration.

Creating arrays

Configuring the arrays involves a few basic steps. All configuration parameters are stored on all hard drives that are members of the specific array. This makes it possible to remove and replace controllers without requiring any configuration changes. The configurations can be viewed or modified through the controller's RS-232 port.

Important

We recommend that you make notes about the arrays you create in case you need to duplicate the configuration for an array. See "Trusting an array" on page 61

Important

When creating arrays and logical drives, make sure that you are in compliance with the following guidelines:

- Maximum drives per array = 12
- Maximum number of arrays = 64
- Maximum number of logical drives = 512

No changes are made until the configuration is saved, so it is possible to quit at any time without affecting the current configuration as long as you do not save those changes. Configuration can be performed while the system is active.

To view an array:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

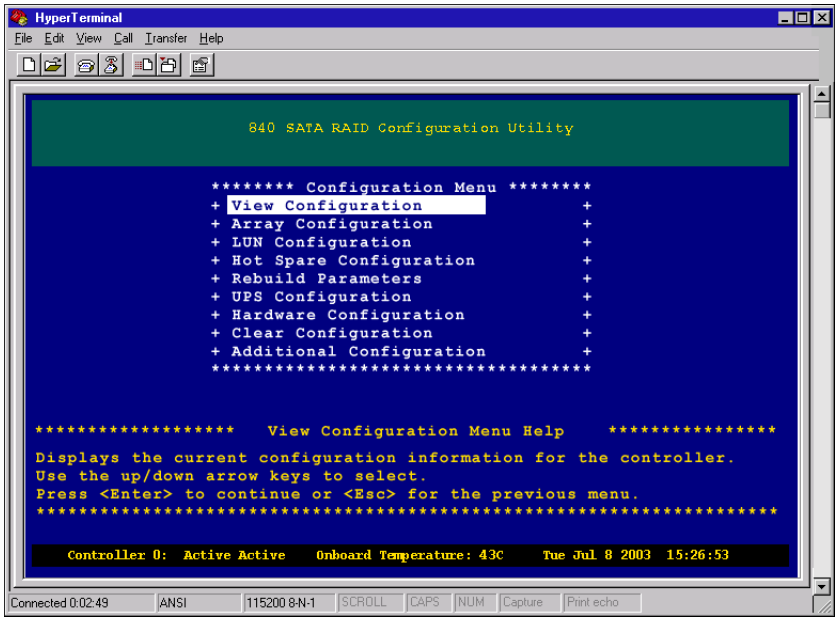

2 From the Configuration Menu, select **View Configuration**, then press ENTER. The *View Configuration Screen* opens.

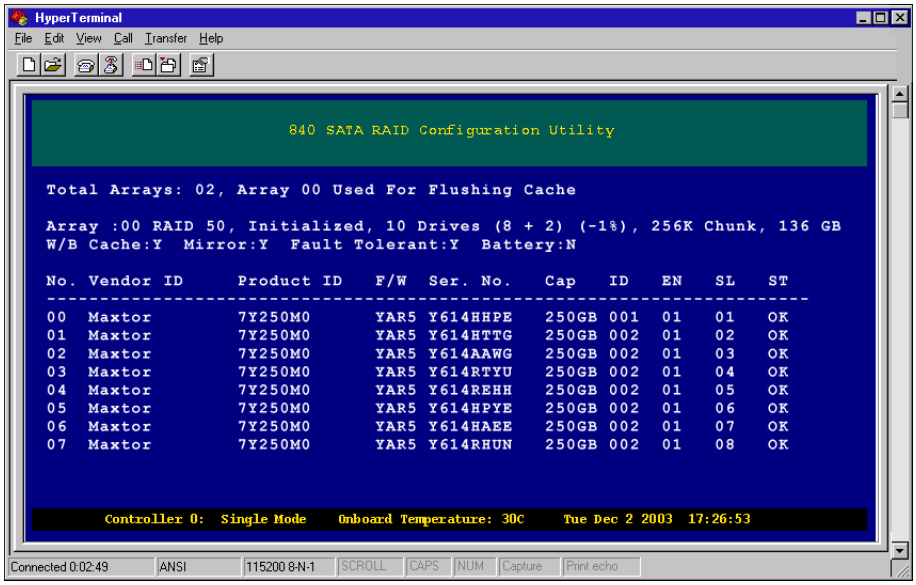

The following table provides a description of the fields shown on the screen.

Field	Description
No.	Drive number in the array.
Vendor ID	Indicates the manufacturer of the drive.
Product ID	Drive model number.
F/W	Drive firmware version.
Ser. No.	Drive serial number.
Cap.	Drive capacity expressed in GBs.
EN, SL	The enclosure and drive slot number, used for drive identification.
ST	The drive status, will either be “OK” or “FL.”

3 Press Esc to return to the previous screen.

To create an array:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

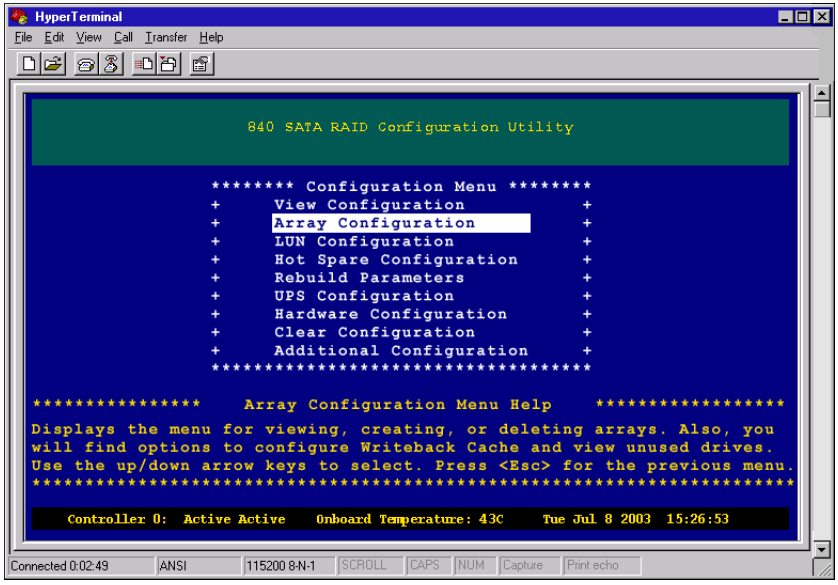

2 To access the Array Configuration Menu, select **Array Configuration,** then press **ENTER**. The *Array Configuration Menu* opens.

The Array Configuration Menu lets you review, create, delete, or otherwise modify the active arrays.

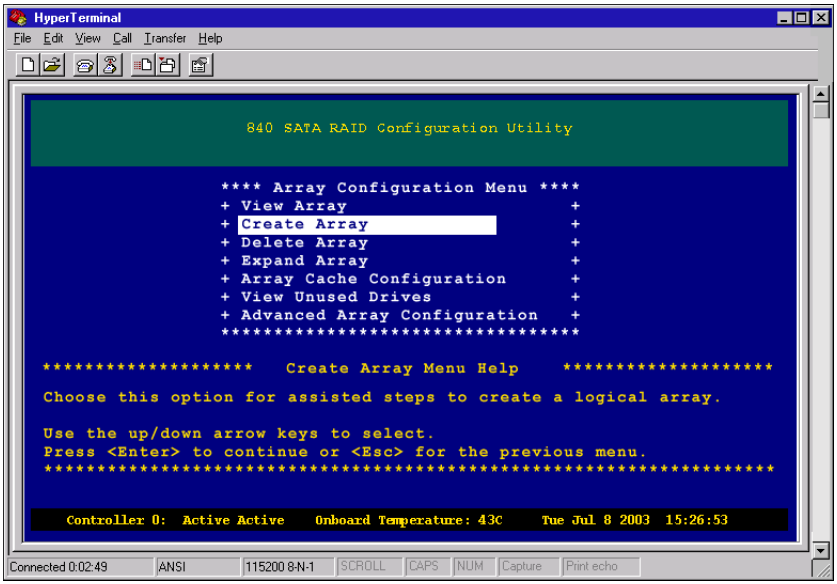

3 From the Array Configuration Menu, select **Create Array**, then press **ENTER**. The *Array Name* screen opens.

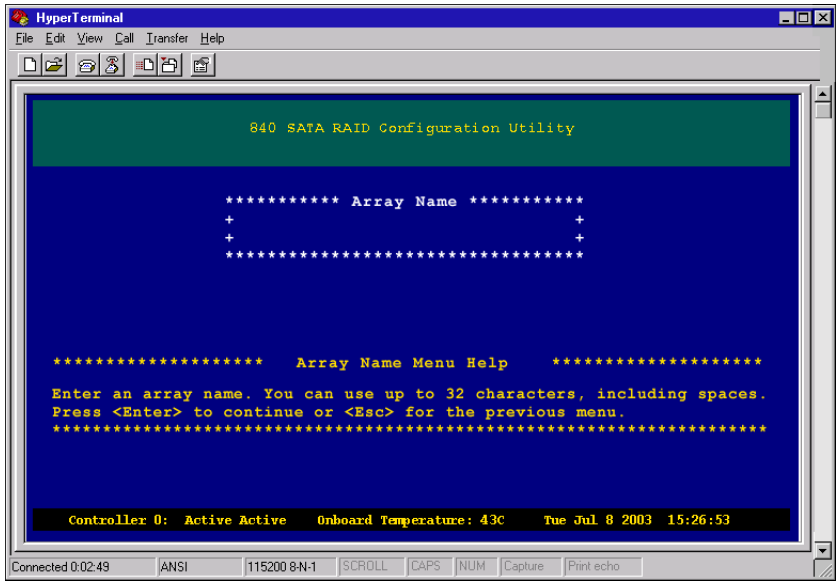

4 Type a name for the array (using as many as 32 characters), then press **ENTER**. The *RAID Level* screen opens.

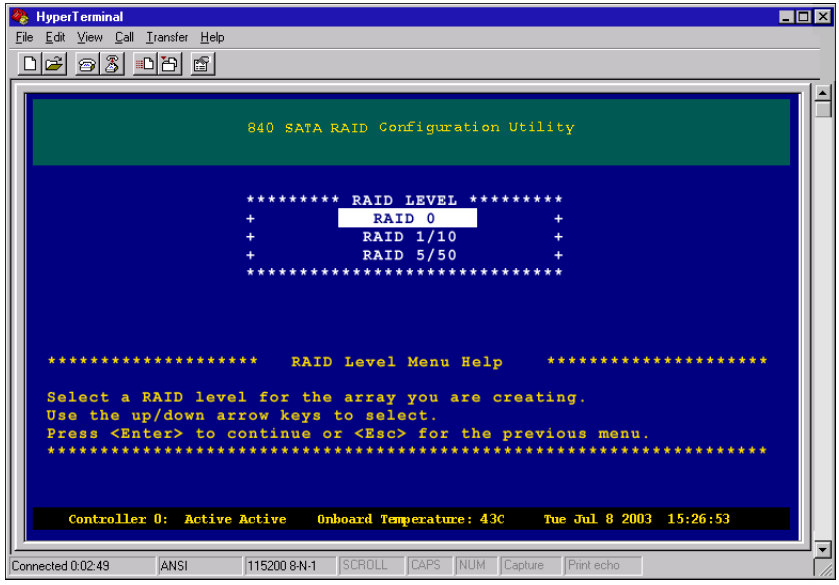

5 Highlight the appropriate RAID level, then press **ENTER**. The *Configuration Type* screen opens.

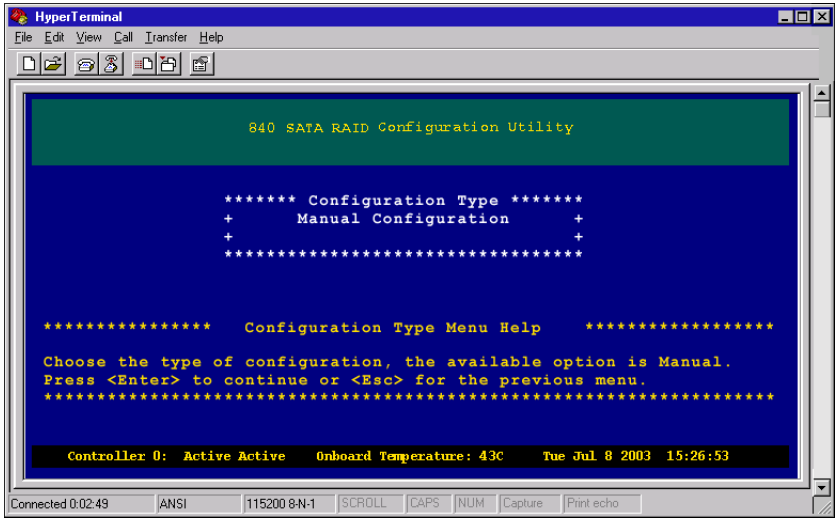

6 Select **Manual Configuration,** then press ENTER. The *Select Chunk Size* screen opens.

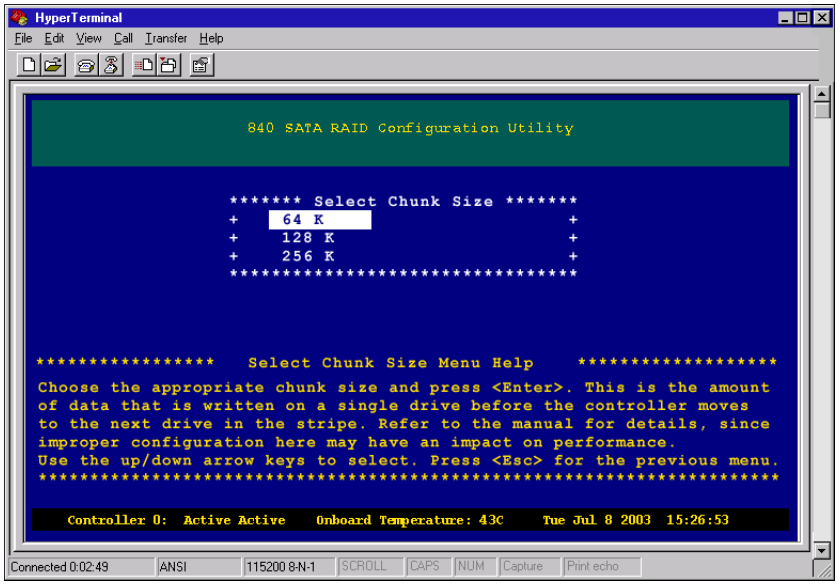

The available chunk sizes are 64 K, 128 K, and 256 K. This is the amount of data that is written on a single drive before the controller moves to the next drive in the stripe.

To achieve optimum RAID 5 write performance you should consider setting the chunk size based on the specified number of drives for a Full Stripe Write when configuring RAID 5/50 arrays. See "Optimizing RAID 5 Write Performance" on page 227 for detailed information.

The primary aim of setting a chunk size is to try to set a stripe size that allows for full stripe writes. The stripe size is determined by the number of data drives multiplied times the chunk size, (8 data drives × 64 K chunk size = 512 strip size).

For maximum performance with RAID 5/50 arrays, you want to do as many full stripe writes as possible. Typically, Windows NT, Windows 2000, and Windows 2003 access at 64 K. Therefore, a stripe size of up to 1 MB would mean the controller has to cluster 16 commands to perform a full stripe write (actually 17 because of alignment). If you were to use a larger stripe size, you run the risk of not being able to cluster sufficiently for the application.

In cases where you are performing larger writes to the controller, you could go up to 2 MB for a stripe size, since you have more data to cluster. **Never exceed 2 MB for a stripe size, since the controller cannot cluster over this size.**

We recommend that you keep the stripe size to 1 MB or less for general use, perhaps increasing it for specific applications. This stripe size is actually the substripe size in RAID 50 cases. A 4+1 array (4 data and 1 parity drives) with a 256 K chunk has a 1 MB stripe size, as does an 8+1 array with 128 K stripe size, and an 8+2, RAID 50 array with 256 K chunk size.

Although 8+1 gives an even stripe size, this does not really matter for an operating system (OS) that writes in significantly smaller chunks. So, 8+1, 10+1, or 11+1 with a chunk size of 64 K would be fine for an OS that does 64 K writes. If the OS does much larger writes, you may want to increase the chunk size. With writeback cache enabled, the controller can keep data and do a full stripe write.

For a single enclosure example - best all around performance:

11+1 with 64 K chunk size would yield a 704 K stripe (11×64=704)

10+2 with 128 K chunk size would yield a 640 K stripe (5×128=640)

When using more drives, 14+2 (64 K), or 12+3 (128 K) should be as good as 12+4. However, we recommend that it be kept to a minimum of 4 data drives for a subarray.

For most sequential accesses, the difference may not be too much for different configurations. However, for large block random writes, it can help if the stripe size is similar to the I/O size. This lets the controller perform a partial full stripe write, where it has most of the data for a full stripe, and can just read some from the drives to complete the stripe.

While Windows NT, Windows 2000, and Windows 2003 do 64 K accesses, these are not aligned. If the chunk size is 64 K, every access crosses a chunk boundary and so involves two read/modify/write commands instead of one. However, since the controller can cluster, this problem is somewhat offset since the controller usually can cluster sufficiently to do full stripe writes. If it is a completely random 64 K access on Window NT, Windows 2000, or Windows 2003, a 64 K chunk is not the best. A chunk size of 128 K or 256 K is better to minimize the number of commands that cross chunk.

Larger chunk sizes should be used if the OS is writing large blocks, or with large sequential writes where the controller can cluster sufficiently. Always keep the stripe size below 2 MB.

7 Select the desired chunk size from the menu, then press **ENTER**. The *Drive Selection Menu* opens.

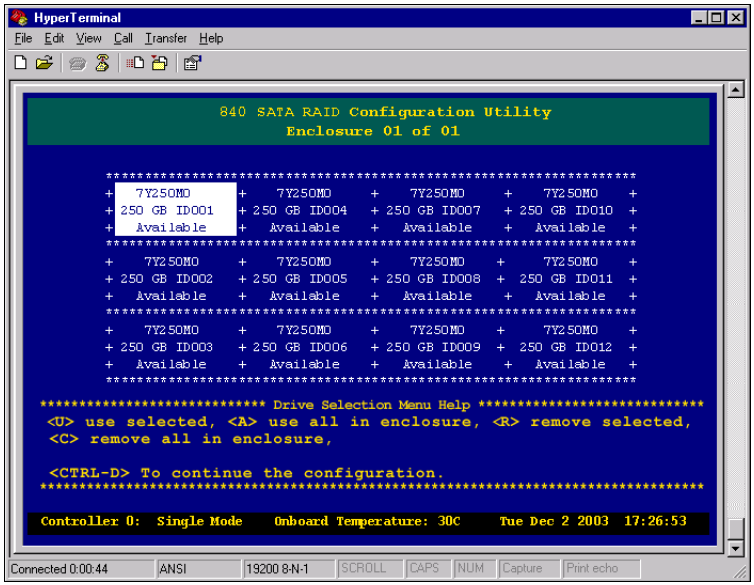

8 Select the drives for the new array by doing the following:

The controller will list all the available drives for each attached enclosure(s). You can choose to add a drive to the array by highlighting the drive, then pressing the **U** key. Or you can select all the available drives by pressing the **A** key. Remove a drive that has been selected by pressing the **R** key, or remove all of the drives selected by pressing the **C** key.

Tips & Tricks

If you have multiple enclosures, the enclosure you are viewing is displayed at the top of the page. To move to the drives in the next enclosure, highlight the last drive in the column or row and press the right or down arrow key.

9 To continue, press the **CTRL** + **D** keys when you have completed your selections. The *Drive Capacity Reserve* screen opens.

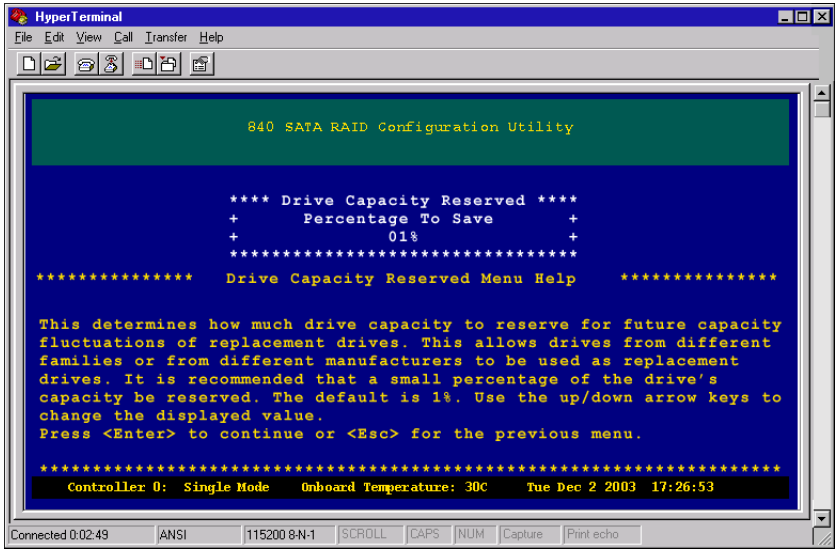

Tips & Tricks

The drive capacity reserve determines how much drive capacity to retain for future capacity fluctuations of replacement drives. (This is not applicable to RAID 0 configurations.) The default is 1%.

10 Type the correct value for the drive capacity reserve (percentage), then press **ENTER**. The *Sub-Array Selection* screen opens.

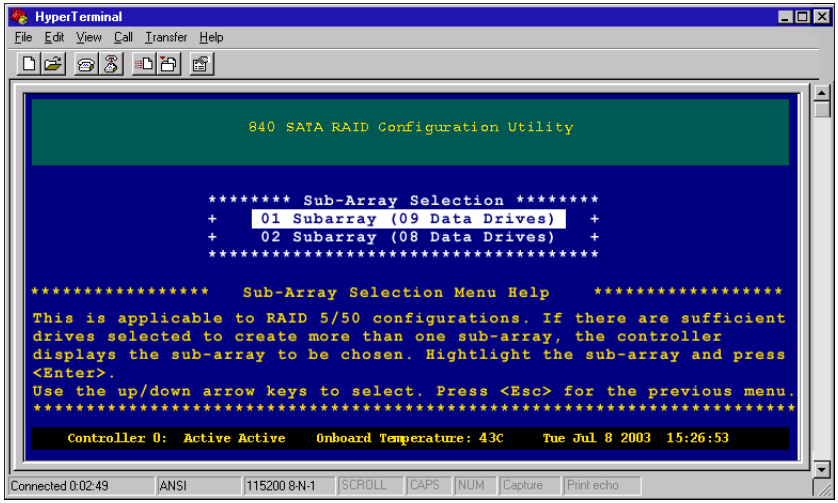

The controller displays the number of available sub-arrays you can choose from. In this example, the first one is a single sub-array that provides nine data drives and one parity drive, and the second example is two sub-arrays that provides two four-drive sub-arrays with one parity drive.

When you create a redundant array, such as a RAID 5 or 50 type, you will always be creating at least one sub-array comprised of the data drives and one parity drive. When a sufficient number of drives have been selected you can choose to make multiple sub-arrays of your array. This provides the benefits of a quicker recovery from a drive failure since only one of the smaller sub-arrays is affected.

Tips & Tricks

If this is just one single array with less than five hard drives, choose the default of one sub-array.

11 Select the number of sub-arrays required, then press **ENTER**. The *Initialize/Trust Arrays* screen opens.

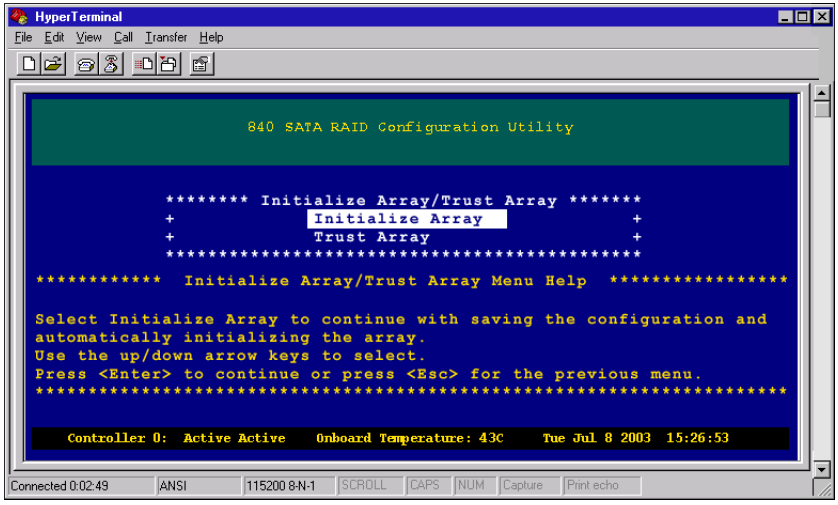

12 Select **Initialize Array**, then press **ENTER**. The *Save Configuration* screen opens.

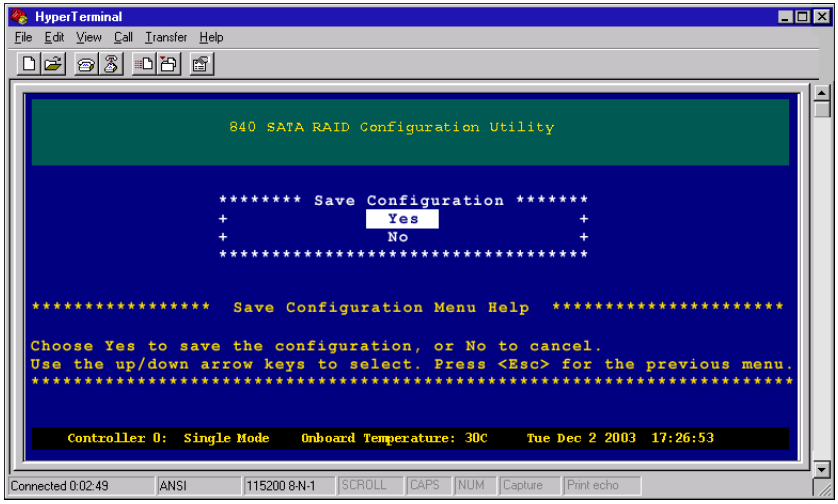

For details on trusting arrays, refer to "Trusting an array" on page 61.

Important

You must initialize the array before using the logical drives created from the array.

13 Save the configuration to create the array. Click **Yes**, then press ENTER. You will see a message that the configuration is being saved.

14 After it has completed the process, press any key to continue.

15 Continue with setting up the Array Read-Ahead Cache and Writeback Cache, then go to the Hot Spare drives options.

Tips & Tricks

After you have completed these tasks, go to the next chapter. Create the logical drives and perform the LUN assignment to complete the setup. Additionally, for more control over the logical drives, see "SAN LUN Mapping" on page 89 and perform your appropriate mappings.

Configuring array read-ahead and writeback cache

The following steps will guide you through configuring the read-ahead cache and writeback cache options for each array. Each array has a different set of cache settings. Read-ahead cache and writeback cache thresholds work together to allow fine tuning of the I/O performance. You will need to set these options for each array you have created.

Read-ahead cache

The Read-Ahead function improves the data retrieval performance by allowing the controller to read into cache a full stripe of data at one time, which greatly improves the cache hits. For smaller transfers the Read-Ahead algorithm can improve performance. If, for example, the stripe size is 256 KB (chunk size multiplied by the number of data drives) and the host requests 64 KB of data, when Read-Ahead is enabled, the controller will read in advance a full 256 KB. When the host request the next 64 KB block, the data will already be in the cache. Depending on the data patterns for your application, disabling the read-ahead cache can help load balance the read and write operations, which can increase performance.

To configure the read-ahead cache:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

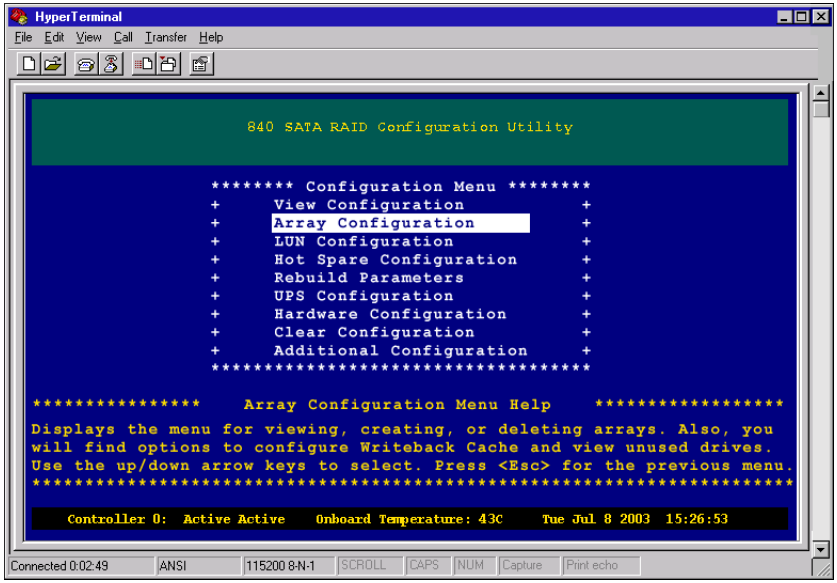

2 From the Configuration Menu, select **Array Configuration**, then press ENTER. The *Array Configuration* Menu opens.

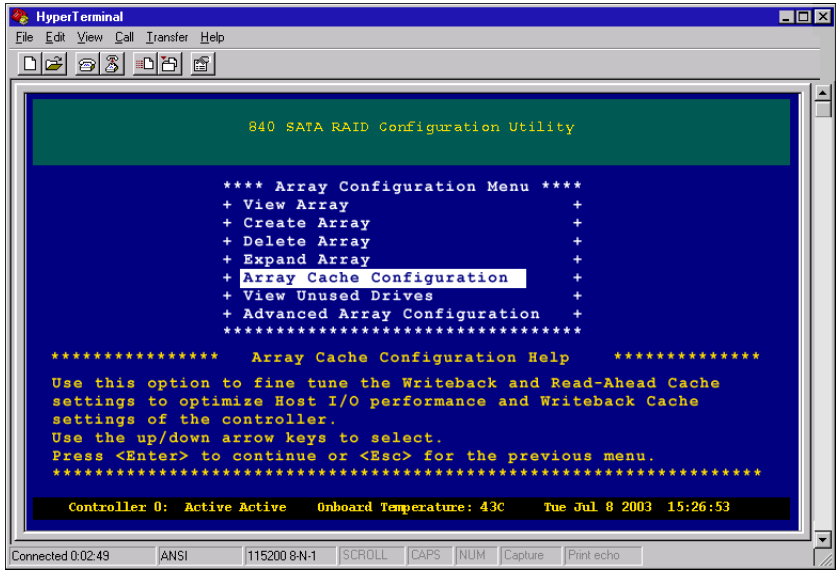

3 From the Array Configuration Menu, select **Array Cache Configuration**, then press ENTER. The *Array Cache Configuration* Menu opens.

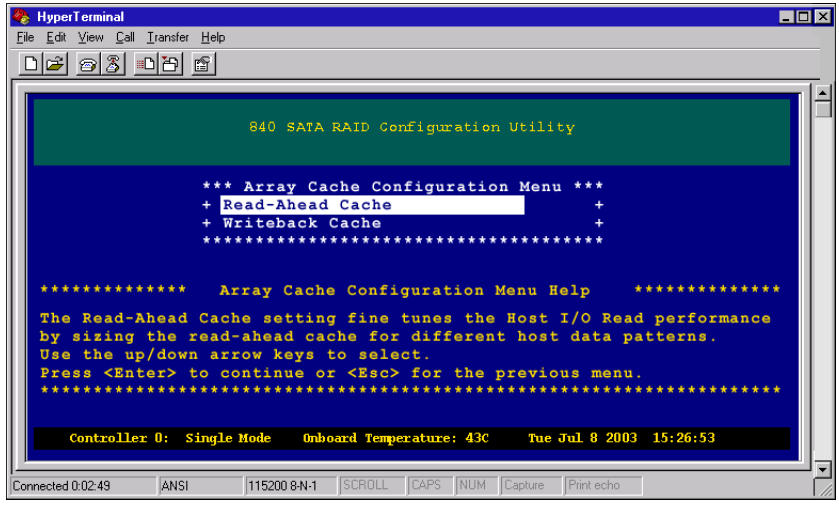

4 From the Array Cache Configuration Menu, select **Read-Ahead Cache**, then press **ENTER**. The *Select Array* screen opens.

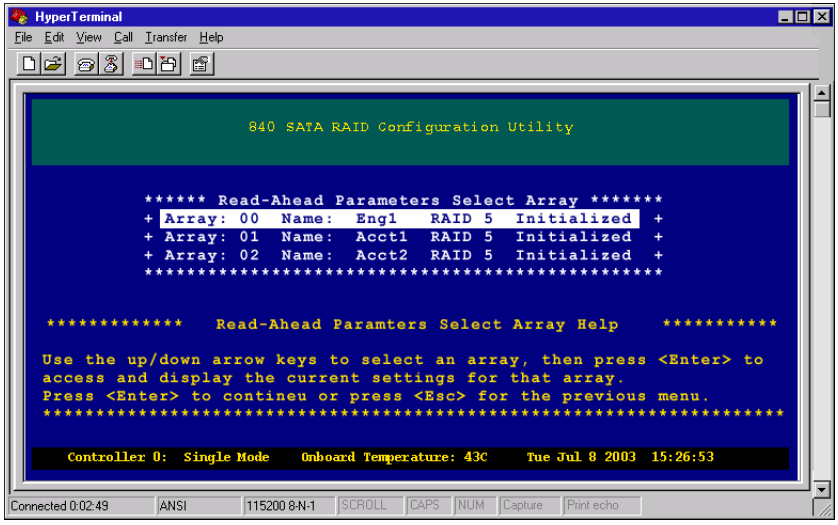

5 Press the up or down arrow keys to select an array, then press **ENTER**.

6 Press the up and down arrow keys to select a Read-Ahead Cache parameter to use for the selected array, then press **CTRL** + **D** to save the changes.

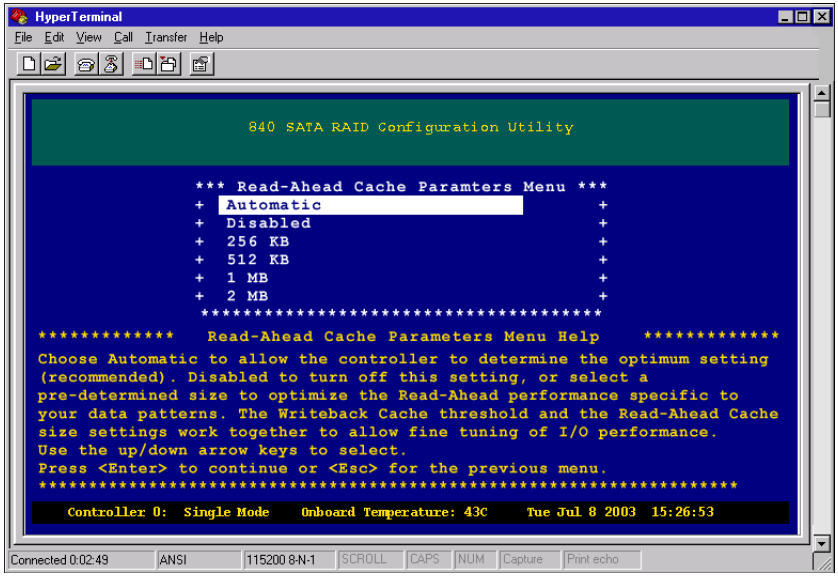

The choices are **Automatic**, **Disabled**, and four pre-determined sizes. Select **Automatic** (the default), and recommended setting. It lets the controller determine the optimum setting. **Disabled** will turn off the read-ahead cache. Choose one of the pre-determined sizes to optimize the read performance based on your data patterns.

7 Save the configuration changes by selecting **Yes**, then press ENTER.

Writeback cache

In a writeback cache operation, data is sent to the controller from the host and the controller immediately responds to the host confirming the data was received and written to the media. The host can then send more data. This can significantly increase performance for host systems that only send a low number of commands at a time. The controller caches the data, and if more sequential data is sent from the host, it can cluster the writes together to increase performance further. If sufficient data is sent to fill a stripe in RAID 5/50 configurations, the controller can perform a full stripe write, which significantly reduces the write overhead associated with RAID 5/50.

Disabling writeback cache ensures that the data is sent to the drives before status is returned to the host. With writeback cache enabled, if a short term power failure occurs, the battery back-up unit provides adequate power to make sure that cache is written to disk when the power is restored. In duplex operations, the cache is mirrored to both controllers which provides further redundancy in the event of a single controller failure.

Mirrored cache is designed for absolute data integrity. The cache in each storage processor contains both primary cached data for the disk groups it owns, and a copy of the primary data of the other storage processor. Mirrored cache ensures that two copies of cache exist on both storage processors, before confirming to the operating system that the write has been completed.

Below is a table from the Gateway 840 Series User Guide, which list the hold-up times for data for the battery backup unit.

Configuration	Memory Vendor and Part Number	Measured Current Draw	Absolute Maximum Backup Time	Expected Safe Backup Time
Main board only w/256 MB	Kingston KVR100X72C2/256	27.9mA	41.2 hours	20.6 hours
Main board w/256 MB and Coprocessor w/256 MB	Kingston KVR100X72C2/256	48.3mA	23.8 hours	11.9 hours

Normally, write-intensive operations will benefit through higher performance when writeback cache is enabled on that array. Read-intensive operations, such as a streaming server, may not benefit from writeback cache.

The writeback cache is used to optimize the write performance specific to your data patterns. In general, larger cache sizes will increase the write performance but can lower simultaneous read performance. The recommended size is 16 MB. The strategy of write operations results in a completion signal being sent to the host operating system as soon as the cache receives the data to be written. The hard drives will receive the data at a more appropriate time in order to increase controller performance.

To configure the writeback cache:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

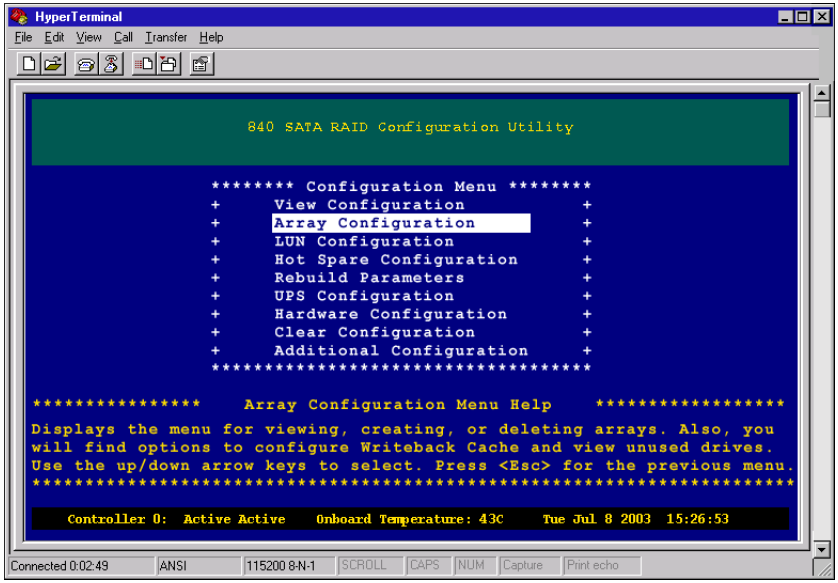

2 From the Configuration Menu, select **Array Configuration**, then press **ENTER**. The *Array Configuration Menu* opens.

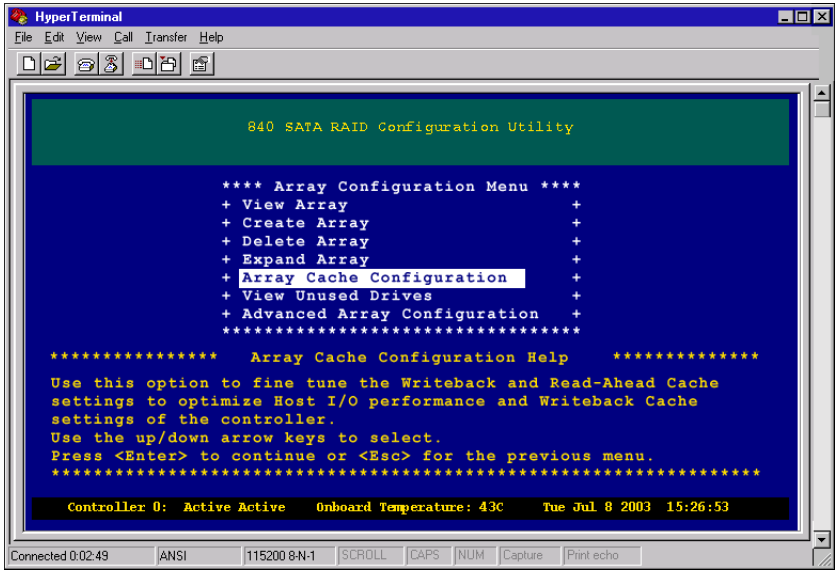

3 From the Array Configuration menu, select **Array Cache Configuration Menu**, then press **ENTER**. The *Array Cache Configuration Menu* opens.

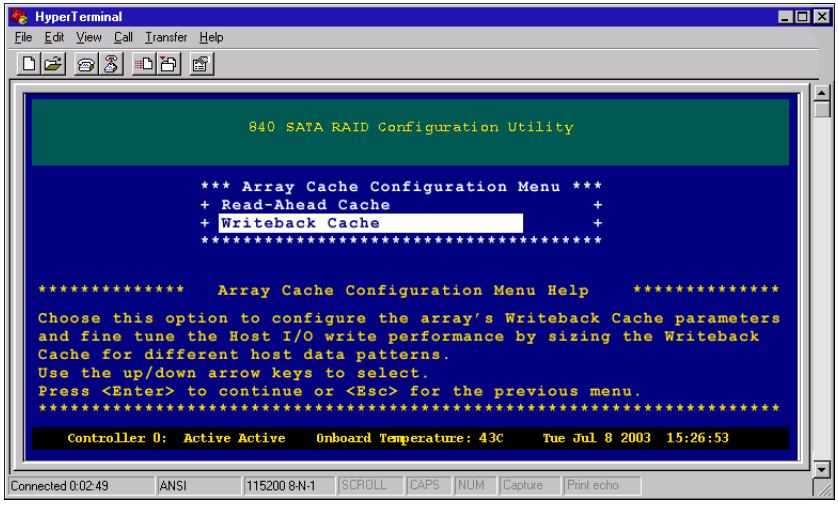

4 From the Array Cache Configuration Menu, select **Writeback Cache**, then press ENTER. The *Writeback Parameters Select Array* menu opens.

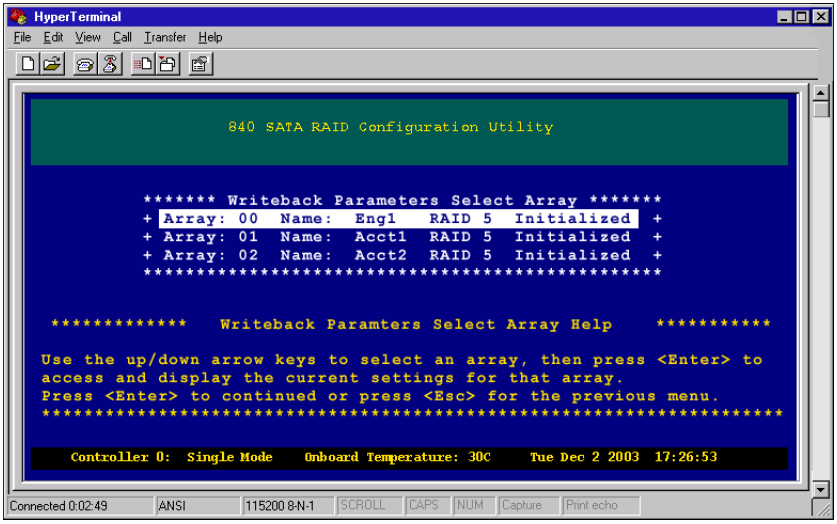

5 Press the up and down arrow keys to select an array from the list to which you want to make changes, then press **ENTER**. The *Writeback Cache Parameters* Menu opens.

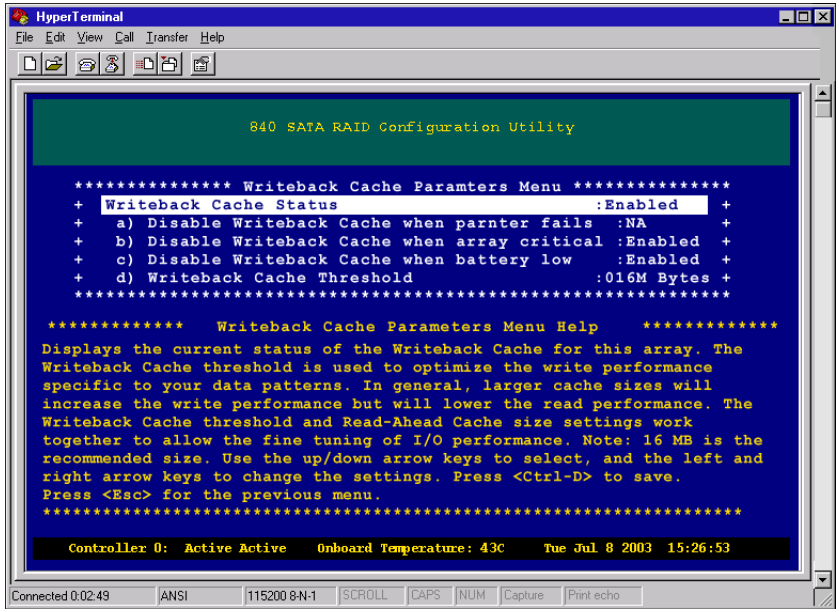

6 Choose a Writeback Cache parameter to use for the selected array.

Press the up and down arrow keys to select the cache, and the left and right arrow keys to toggle from Disabled to Enabled, or choose one of the pre-determined cache threshold sizes (1 MB, 2 MB, 4 MB, 8 MB, 16 MB, 32 MB, 64 MB, 128 MB, or 256 MB).

There are three additional options to an active Write Back Cache: Disable if a partner controller fails or is missing, Disable if a controller battery fails, or Disable if the array is in a critical state, (for example, during a rebuild). Enable the options for your application. For maximum data protection, we recommend that you enable all three options.

7 Press **CTRL** + **D** to save the changes. The *Save Configuration* screen opens.

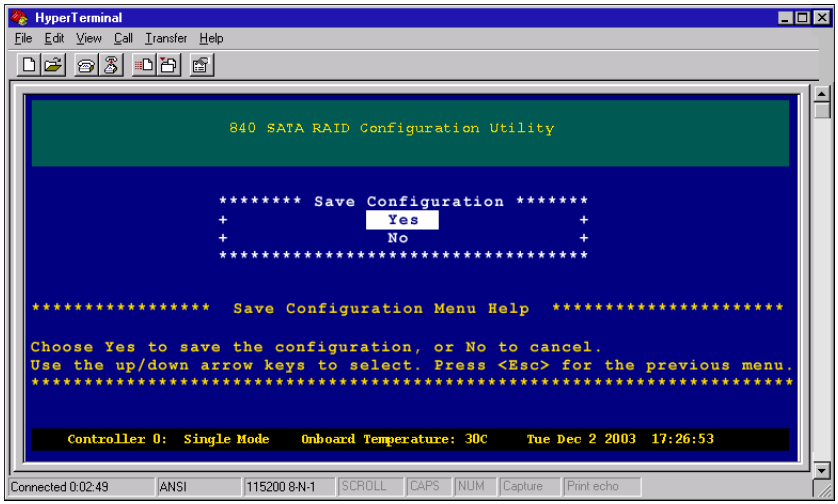

8 Select **Yes** to save the configuration changes, then press **ENTER**.

9 Press any key to continue.

Assigning hot spare drives

The process of configuring fault tolerant arrays includes assigning drives for global or dedicated hot spares. In the event of a drive failure, the controller will use a global hot spare to replace the failed drive in any array. If a dedicated spare is assigned to the specific array, that array will use its dedicated hot spare.

This step is accomplished through the Hot Spare Configuration menu.

To configure hot spare drives:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

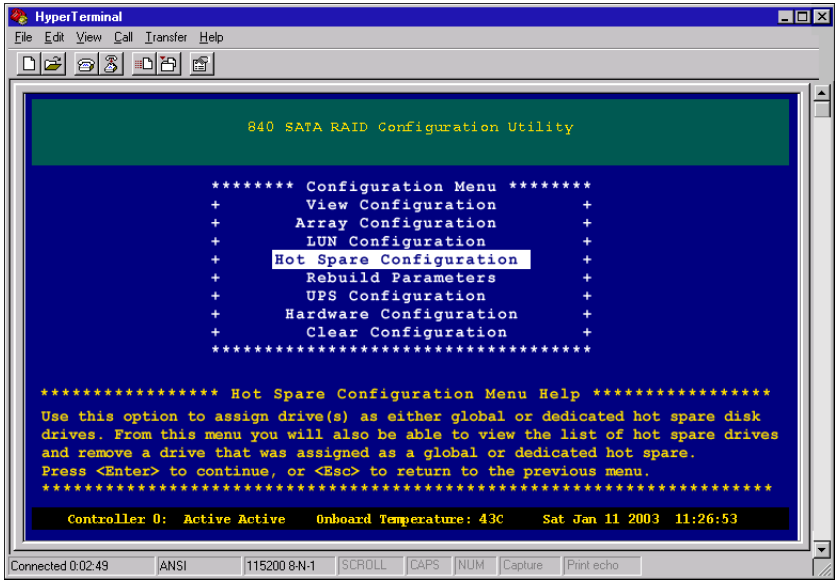

2 From the Configuration menu, press the up or down arrow keys to select *Hot Spare Configuration* menu, then press **ENTER**. The *Hot Spare Configuration Menu* opens.

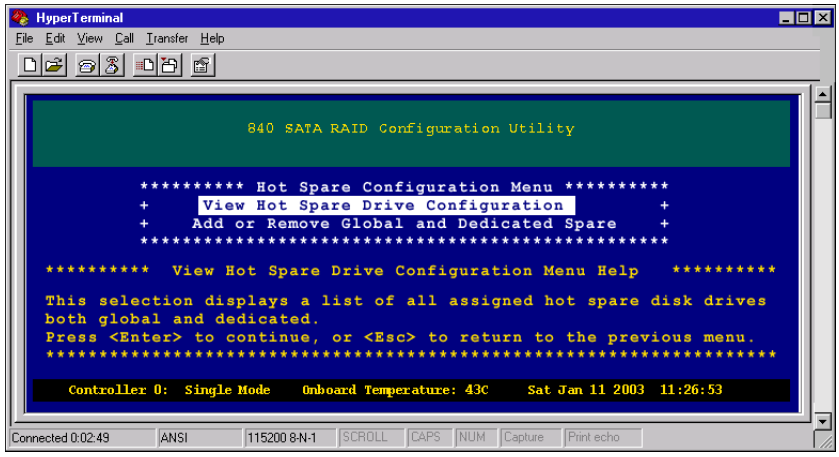

3 To add or remove a hot spare drive, select **Add or Remove Pool and Dedicated Spare**, then press **ENTER**. The *Add or Remove Global and Dedicated Spares* screen opens.

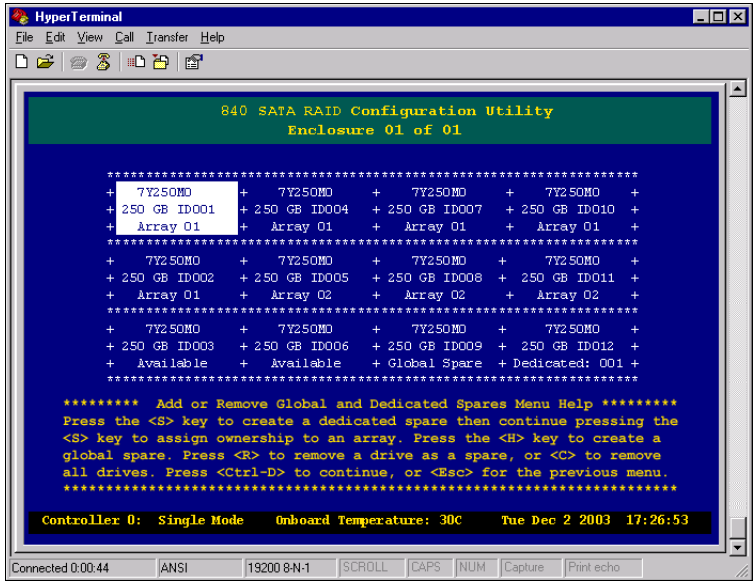

4 Press the arrow keys to highlight an *available* drive to use, then press **S** to assign that drive as a dedicated spare. Continue to press **S** to display the arrays from which you can assign the dedicated spare.

   To add a drive as a pool spare (global), press the arrow keys to highlight an *available* drive and press **H**.

   To remove a drive as a hot spare, highlight the subject drive and press **R**, or press **C** to remove all drives in the enclosure that are currently assigned as spares.

5 After you make your selection, press **CTRL** + **D** to continue.

6 When you are prompted to save the changes, press **Y** to save the new hot spare configurations.

7 Press any key to continue, then go to the next chapter, "Logical Drives" on page 63.

Deleting an array

To delete an array:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

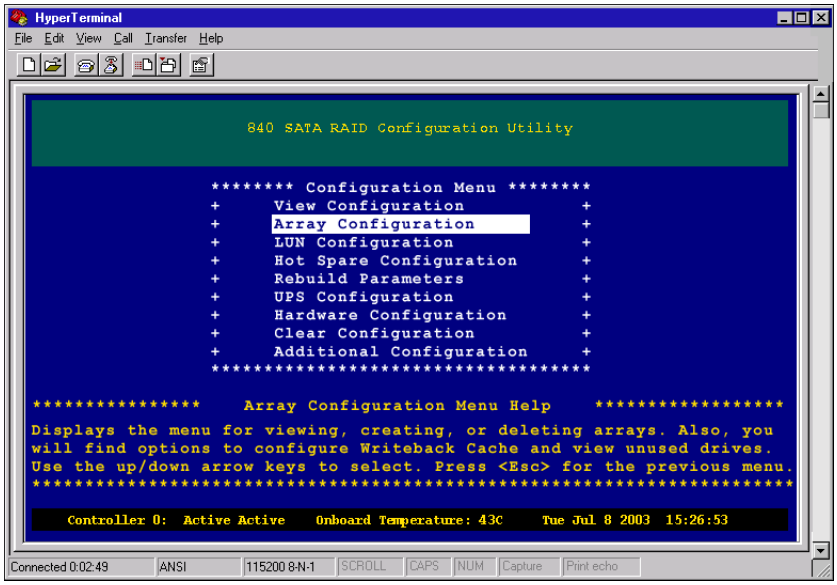

2 From the Configuration Menu, select **Array Configuration Menu**, then press **ENTER**. The *Array Configuration Menu* opens.

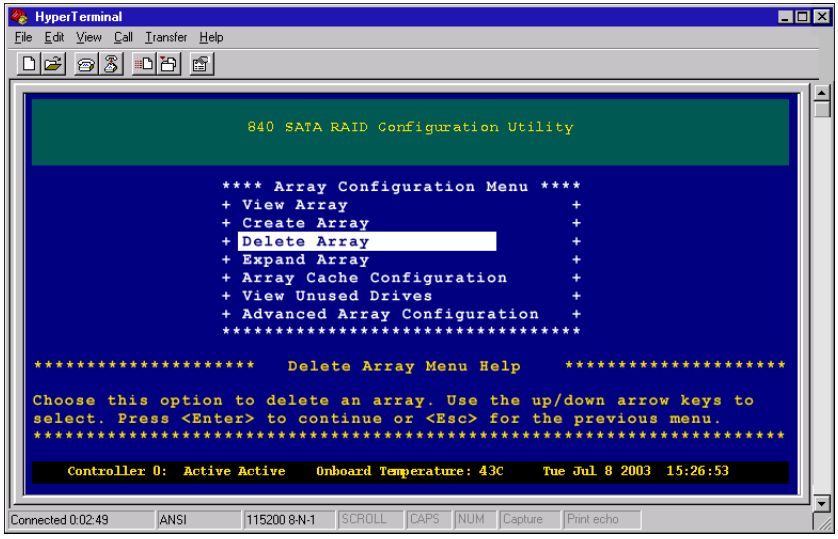

3 From the Array Configuration Menu, select **Delete Array,** then press **ENTER**. The *Delete Array* screen opens.

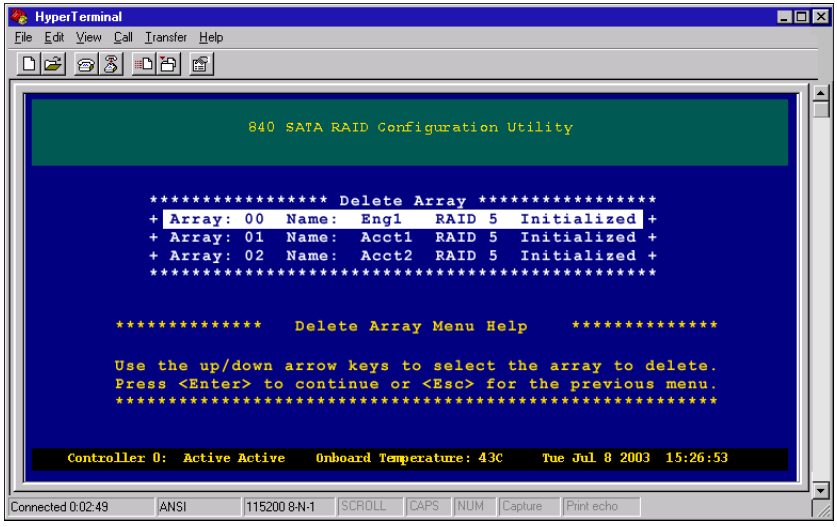

4 Press the up and down arrow keys to select the array to delete, then press ENTER.

5 To complete the deletion, select **Yes**, then press ENTER.

6 Press any key to continue.

Expanding an array

The Expand Array option lets you expand the capacity of your existing array by adding more drives or adding sub-arrays.

Using a RAID 5 array example, adding more drives to the array lets you increase the capacity of that array. If you add additional sub-arrays it requires the exact number of drives to be available for the original sub-array. For example, if you have an array composed of 4 data drives and 1 parity drive, you must have a minimum of five drives available to be able to add at least 1 sub-array.

When you view the array configuration information, the RAID type and number of drives are displayed. These drives are depicted as data drives plus parity drives. In this example, it would be displayed as “4 Drives (3 +1)” which indicates the array/sub-array is composed of 4 drives of which 3 are data drives and 1 is a parity drive. If you were to add 1 sub-array to this particular configuration, the resulting display would be “8 Drives (6+2)” in which case you now have 6 data drives and 2 parity drives. Your RAID 5 array now becomes a RAID 50 array because it’s composed of two sub-arrays, each sub-array has three data drives and one parity drive.

If you are expanding a RAID 50 array by adding drives to the array, you must add an equal number of drives to each sub-array.

The process of expanding the array by adding additional sub-arrays allows for quicker recoveries in the event of a drive failure in any one of the sub-arrays, and the time required to rebuild the array is significantly shortened.

To add more drives to your existing arrays, refer to “Adding additional drives” on page 54.

To add more sub-arrays to your existing array, refer to “Adding additional sub-arrays” on page 57.

Adding additional drives

To add additional drives:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

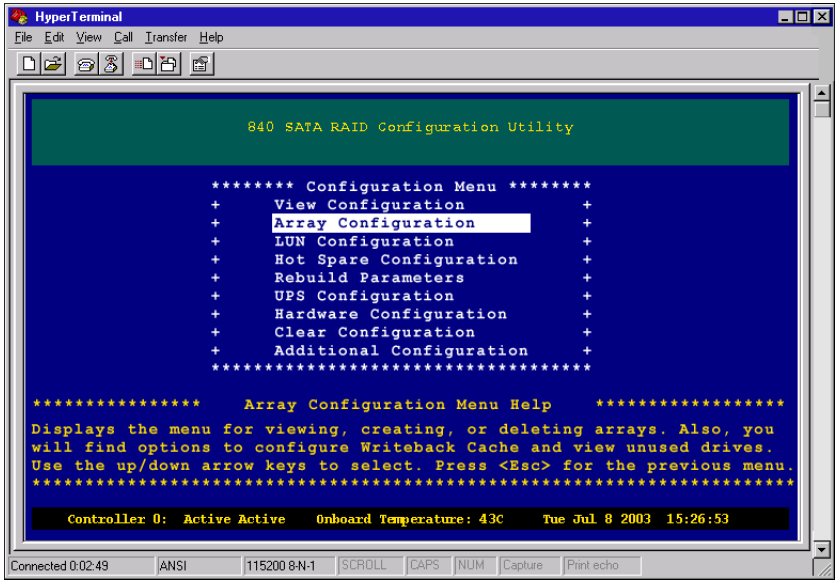

2 From the Configuration Menu, select **Array Configuration Menu**, then press **ENTER**. The *Array Configuration Menu* opens.

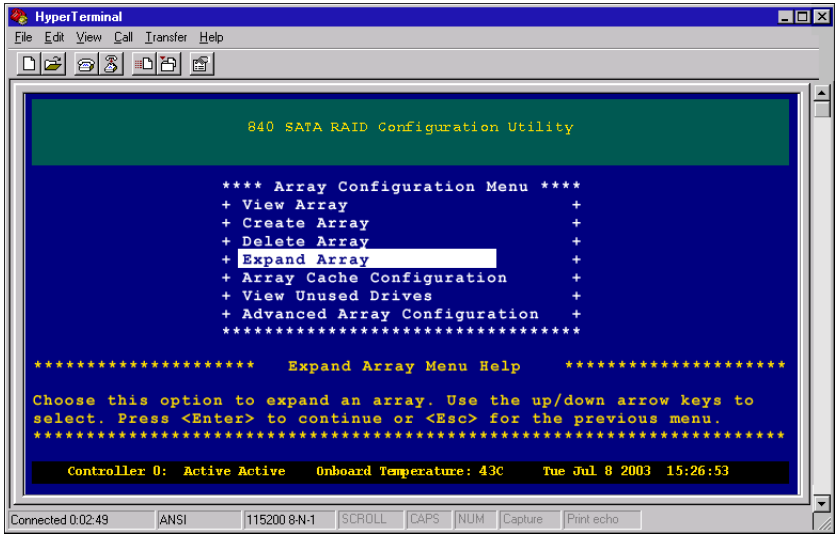

3 From the Array Configuration Menu, select **Expand Array**, then press **ENTER**. The *Select Array to Expand Menu* opens.

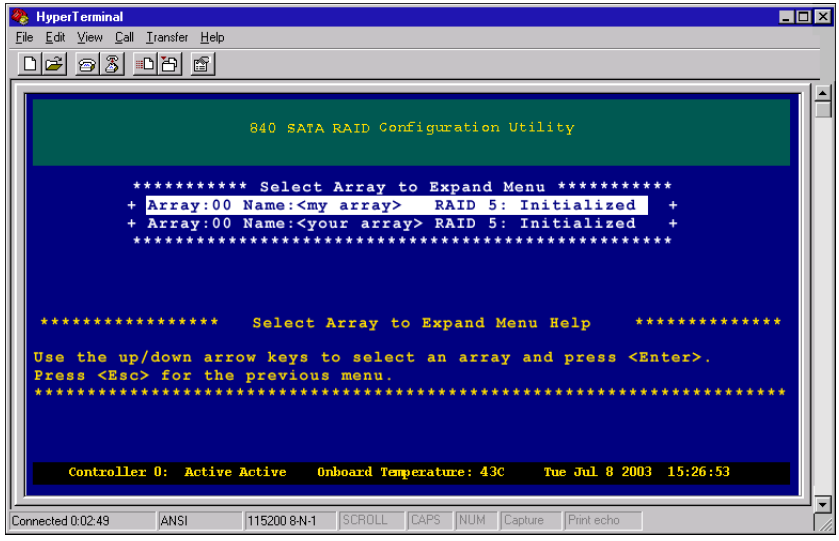

4 Press the up and down arrow keys to select an array to expand, then press **ENTER**. The *Expansion* menu opens.

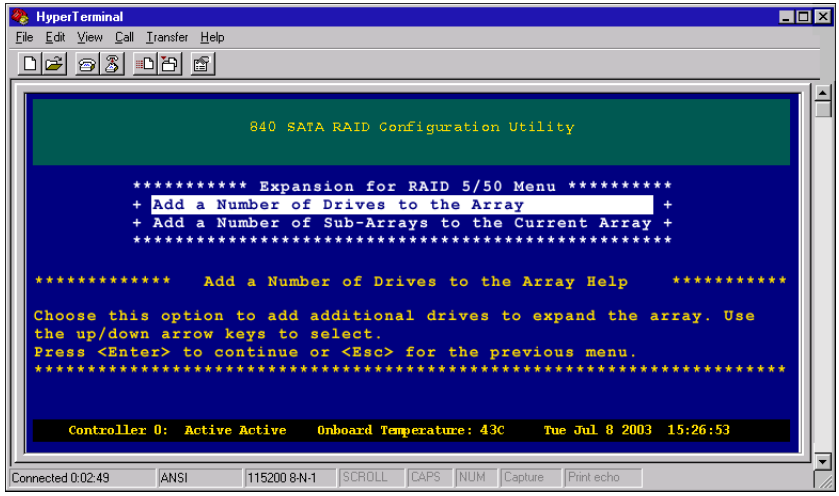

5 From the Expansion menu, select **Add a Number of Drives to the Array**, then press **ENTER**. The *Expand Array Drive Selection* screen opens.

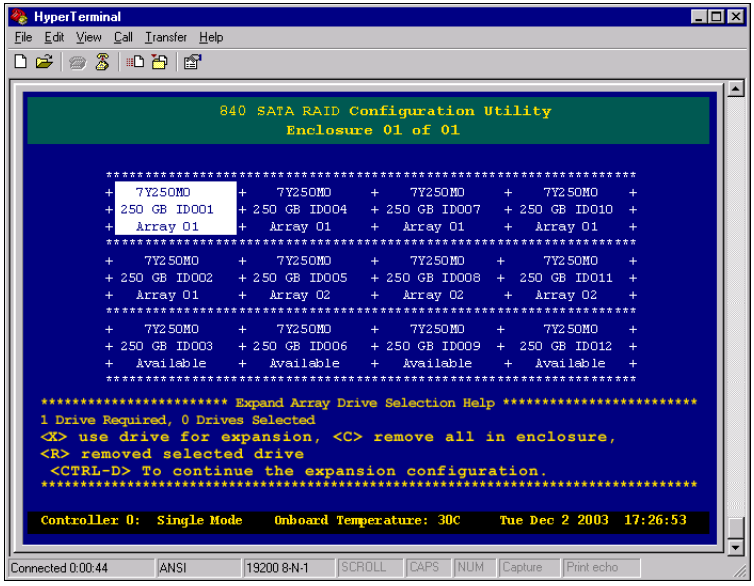

6 Press the up and down arrow keys to select a drive, then press **X** to mark that drive for expansion. Continue to select additional drives if necessary, otherwise press **CTRL** + **D** to continue.

You can cancel your selection by highlighting the drive and pressing **R**, or remove all selected drives and start again by pressing **C**.

7 When you are prompted to save the changes, select **Yes**, then press **ENTER**.

8 Press any key to continue.

Adding additional sub-arrays

To add additional sub-arrays:

1 From the Main Menu, select **Configuration Menu**, then press **ENTER**. The *Configuration Menu* opens.

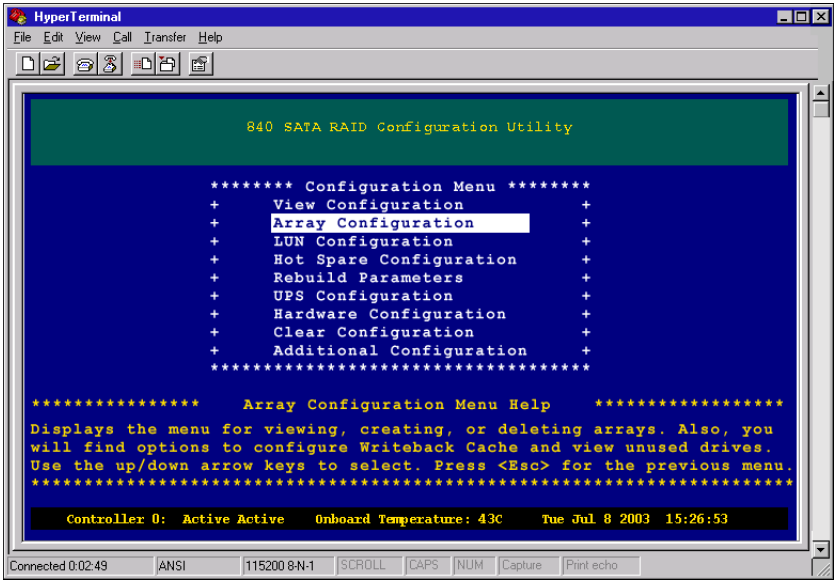

2 From the Configuration menu, select **Array Configuration Menu**, then press **ENTER**. The *Array Configuration Menu* opens.

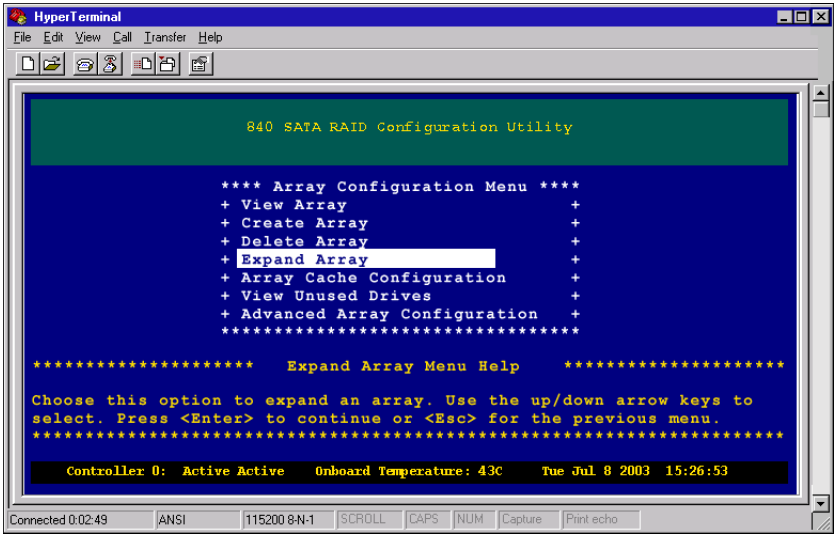

3 From the Array Configuration Menu, select **Expand Array,** then press **ENTER**. The *Select Array to Expand Menu* opens.

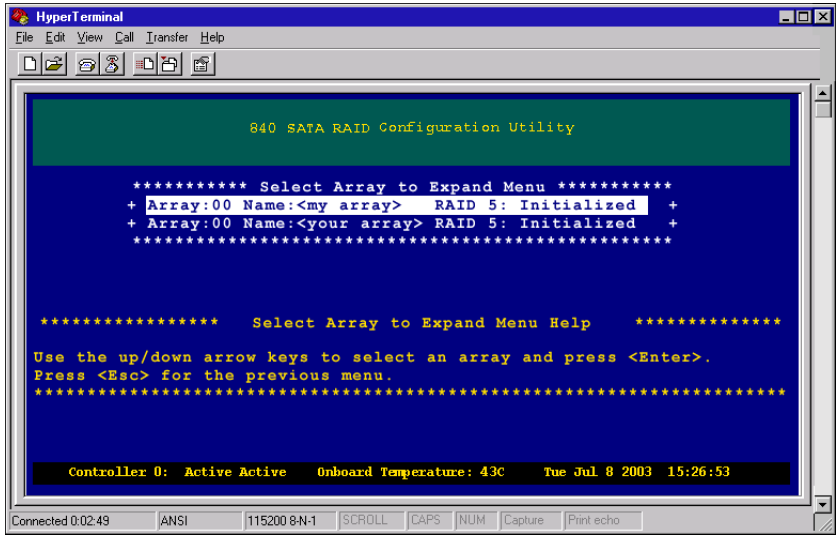

4 Press the up and down arrow keys to select an array to expand, then press **ENTER**. The *Expansion Menu* opens.

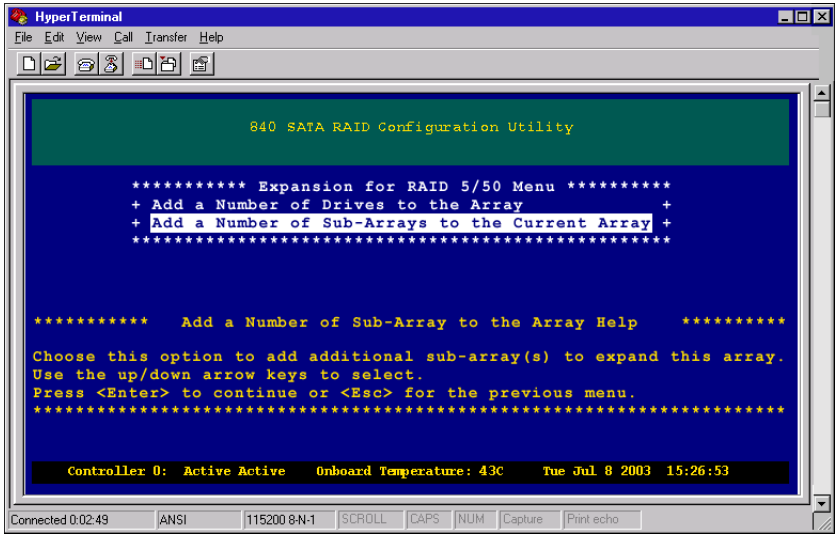

5 Press the up and down arrow keys to select **Add a Number of Sub-Arrays to the Current Array**, then press **ENTER**. The *Expand Array Drive Selection* screen opens.

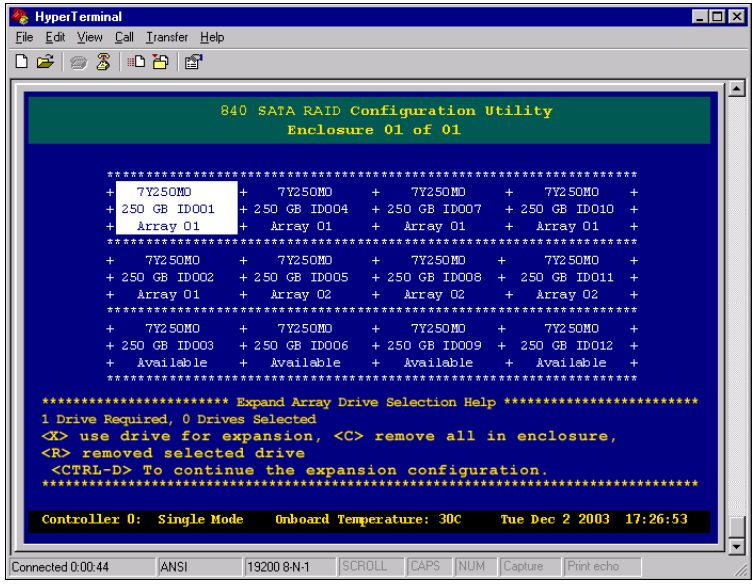

6 Press the up and down arrow keys to select a drive, then press **X** to mark that drive for expansion. Continue to select additional drives if necessary, otherwise press **CTRL** + **D** to continue.

You can cancel your selection by highlighting the drive and pressing **R**, or remove all selected drives and start again by pressing **C**.

7 When you are prompted to save the changes, select **Yes**, then press **ENTER**.

8 Press any key to continue.

Trusting an array

When you create a RAID 5 or 50 array, you have the option to trust the array. This option should *only* be used in environments where you fully understand the consequences of the function. Trust array option is provided to allow immediate access to an array for **testing** application purposes only.

Trust array does not calculate parity across all drives and therefore there is no known state on the drives. As data is received from the host, parity is calculated as normal, but it occurs on a block basis. There is no way to guarantee that parity has been calculated across the entire drive. The parity data will be inconsistent, so a drive failure within a trusted array will cause data loss.

To trust an array:

1. When creating a RAID 5/50 array, after you have selected the number of sub-arrays, the *Initialize Array/Trust Array* menu opens, where you can Initialize or Trust the array.

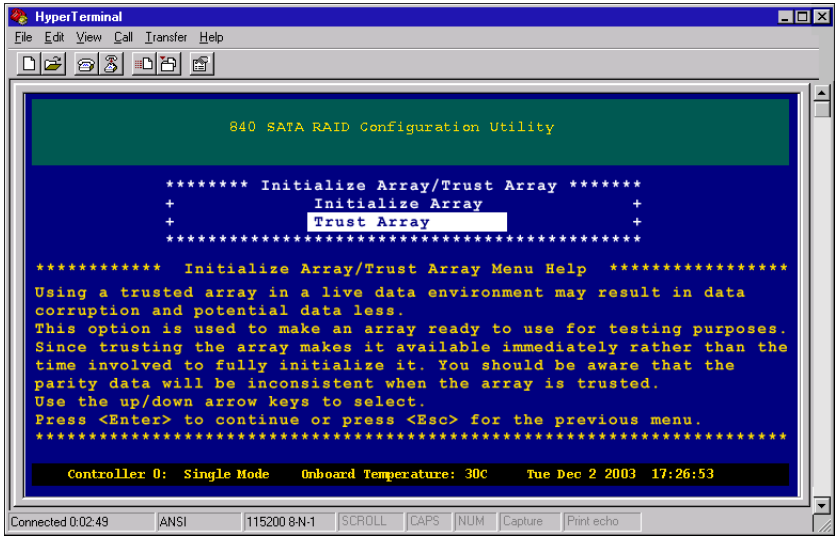

2. (Test Purposes Only) Select **Trust Array,** then press ENTER.

3 You are prompted to save the configuration. Select *Yes* to save the configuration, then press **ENTER**.

4 After the process has completed, press any key to continue.

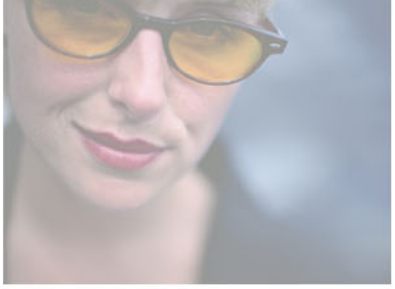

4 Logical Drives

This chapter provides information on creating logical drives. Read this chapter to learn how to:

- Configure a logical drive
- Access the LUN configuration menu
- View unassigned free space
- Create a logical drive

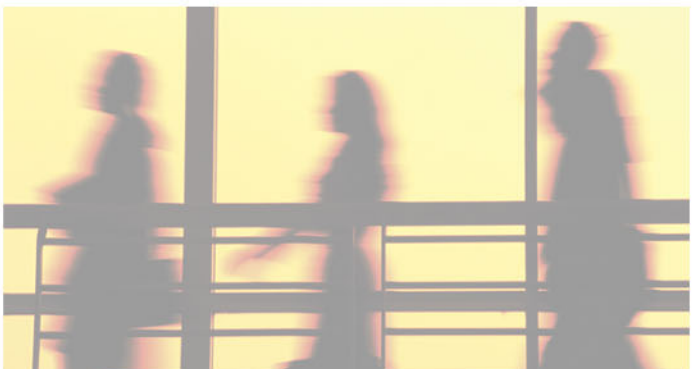

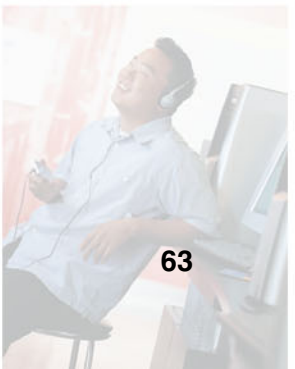

Overview

A *Logical Drive* is defined as a region or combination of regions of unused space on the array(s) which makes the logical drives available to the host operating systems as a disk. You can create up to 512 logical drives. After an array has been created, this region is first marked as unassigned. One or more logical drives can be created in this region or an existing logical drive can be expanded using this region.

A logical drive can be created or expanded in 1 GB increments with a maximum total size per drive of 2,198 GBs. This corresponds to the SCSI 32-bit addressing limitation of 2 TB.

Important

Before you create more than one logical drive, you must be sure that your host HBA and host operating system is setup to handle the desired number of logical drives (LUNs or Logical Unit Numbers). If your operating system does not support multiple logical drives, the host will only be able to see the first logical drive.

Important

For Microsoft® Windows® NT there is a limitation of 231 logical drives. A hot fix is available from Microsoft. See Microsoft Knowledge Base Article-245637.

Terminology

The following table describes the terminology relating to logical drives.

Term	Description
Segmentation	Any logical drive can be expanded into any free region, so it is possible to easily add capacity at any time. There is no requirement that any additional space be contiguous. Logical drive segmentation is completely transparent to the host systems.
Availability	To accommodate hosts with multiple ports and multiple host systems, you can restrict a logical drive's availability to a particular controller port. Access can be enabled or disabled for each port of each controller.

Mapped Logical Drive Number	Each logical drive is presented to the host system with a unique LUN. In certain cases (such as after deleting another logical drive) it may be desirable to change the number of the logical drive. This can be done at any time, bearing in mind that any attached host systems may need to be rebooted or reconfigured to retain access.
Unassigned Free Space	The controller keeps a map of all the space that is not assigned to any logical drive. This space is available for logical drive creation or expansion. Each unassigned region is individually listed.

Accessing the LUN configuration menu

To access the LUN configuration menu:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

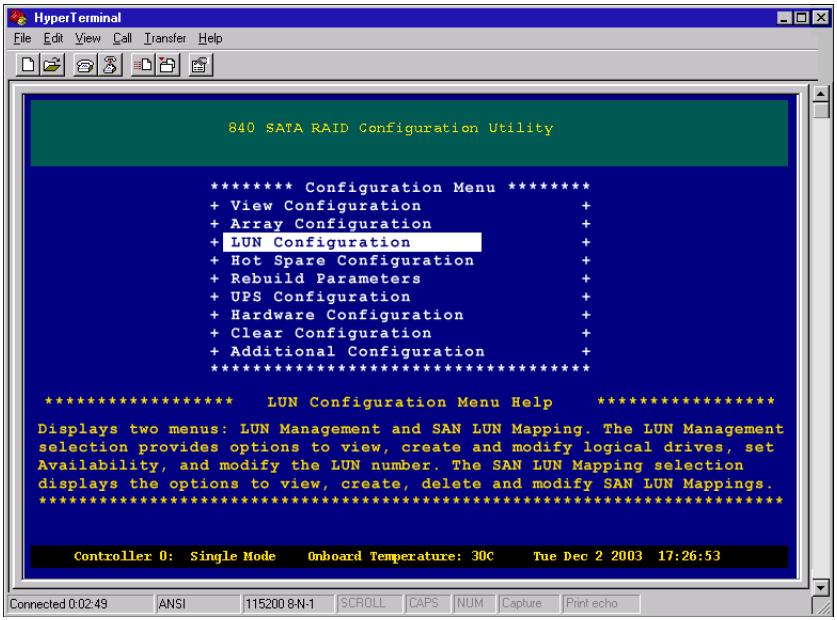

2 From the Configuration Menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

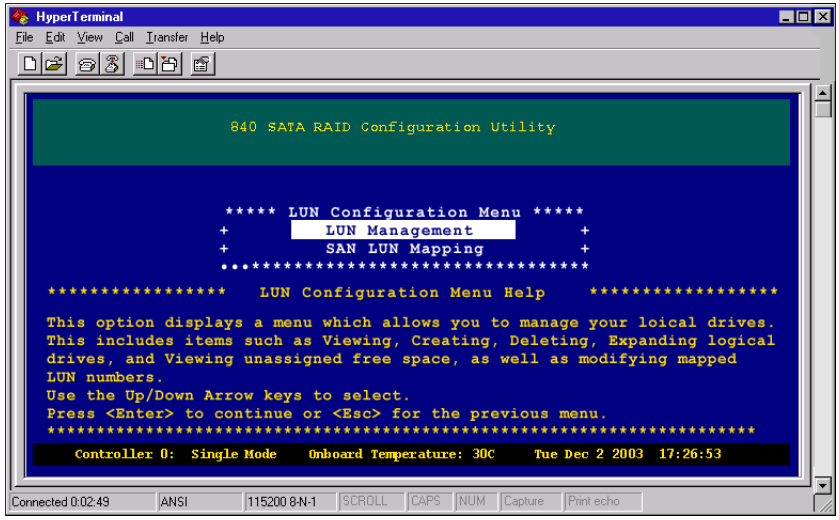

The *LUN Configuration Menu* has two menu options: one for managing logical drives, which includes utilities for viewing, creating, deleting, and managing LUNs, and the other is for performing SAN LUN mapping operations. For more information see "SAN LUN Mapping" on page 89

Viewing unassigned free space

Prior to creating a logical drive, you may want to review the available unassigned free space. This will help you to identify the unused regions or segments for use during the creation of your logical drives. Normally with a first time configuration this is not necessary, since all of the space is unassigned.

To view unassigned free space:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

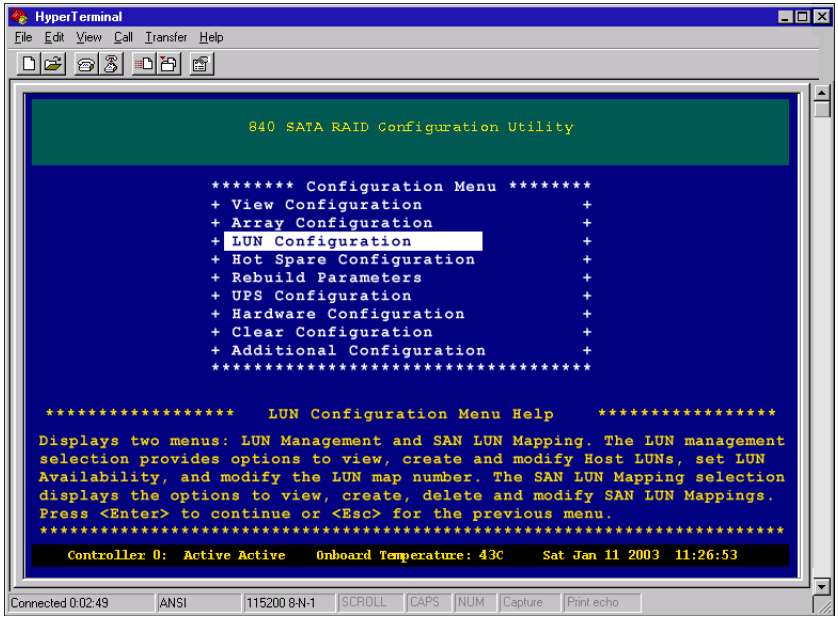

2 From the Configuration Menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

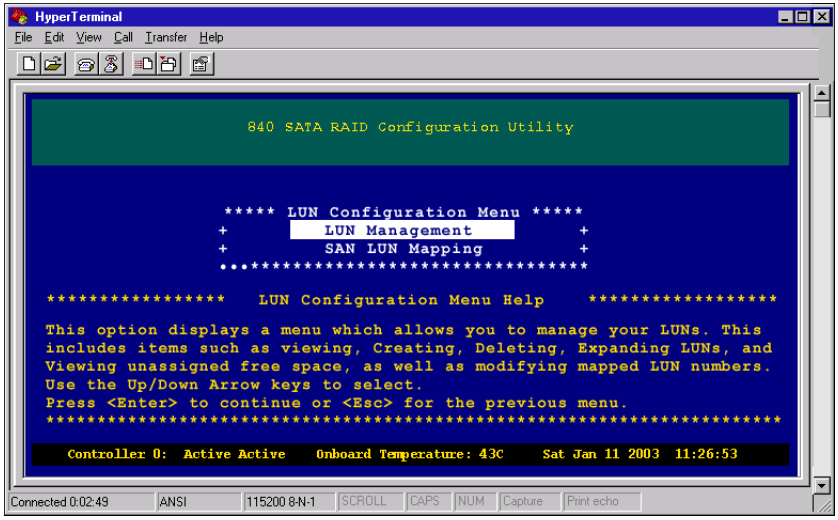

3 From the LUN Configuration Menu, select **LUN Management**, then press **ENTER**. The *LUN Management Menu* opens.

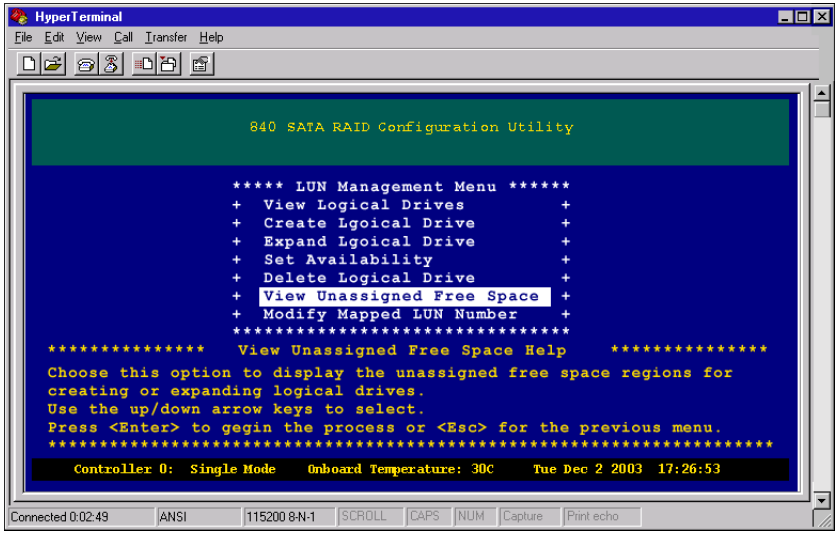

4 From the LUN Management Menu, select **View Unassigned Free Space**, then press **ENTER**. The *Unassigned Free Space* screen opens.

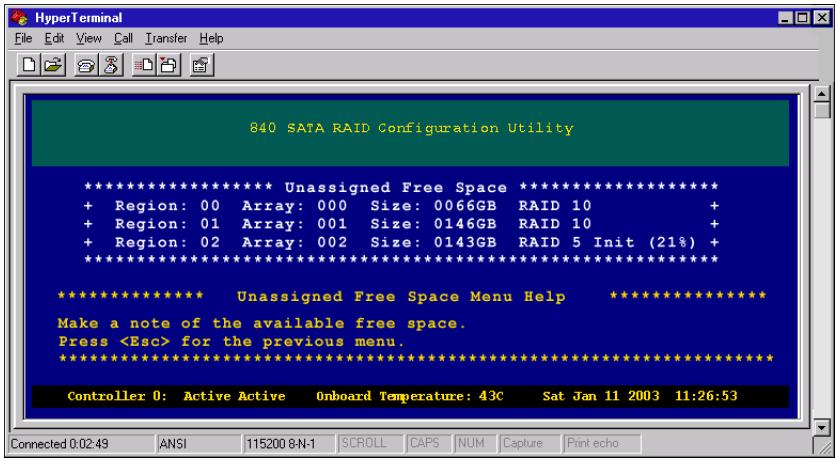

5 Review the free space regions and note them for later use.

Creating a logical drive

To complete the process of configuring your arrays, you will need to create one or more logical drives. Creating a logical drive from the available free space regions presents the logical drive to the host operating system.

To create a logical drive:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

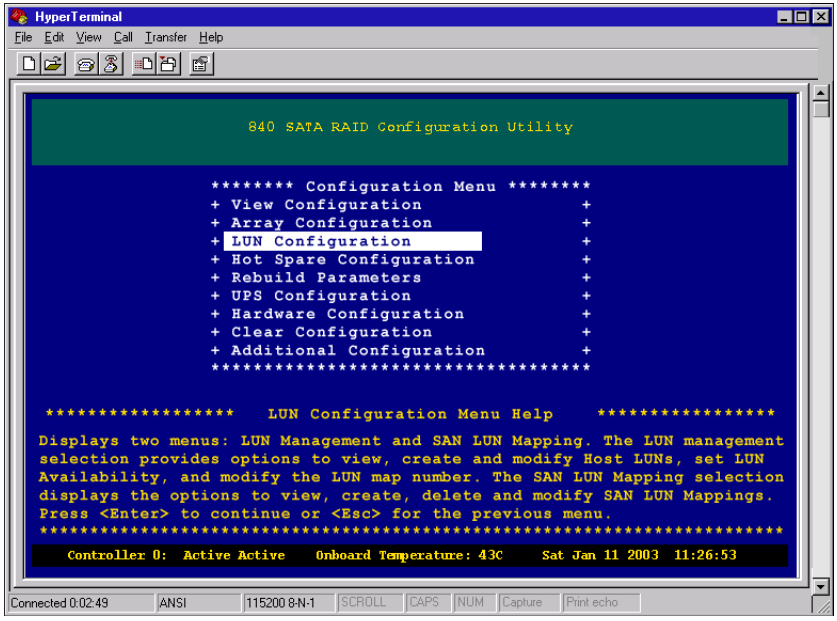

2 From the Configuration menu, select **LUN Configuration**, then press ENTER. The *LUN Configuration Menu* opens.

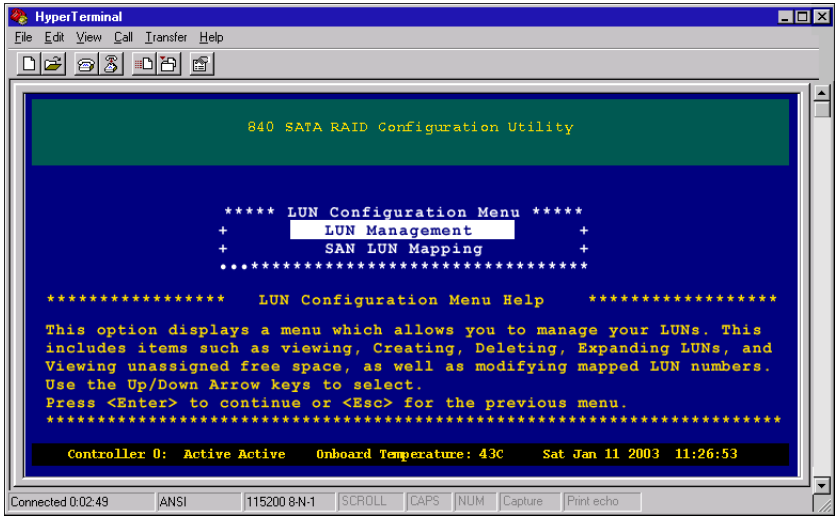

3 From the LUN Configuration menu, select **LUN Management**, then press ENTER. The *LUN Management Menu* opens.

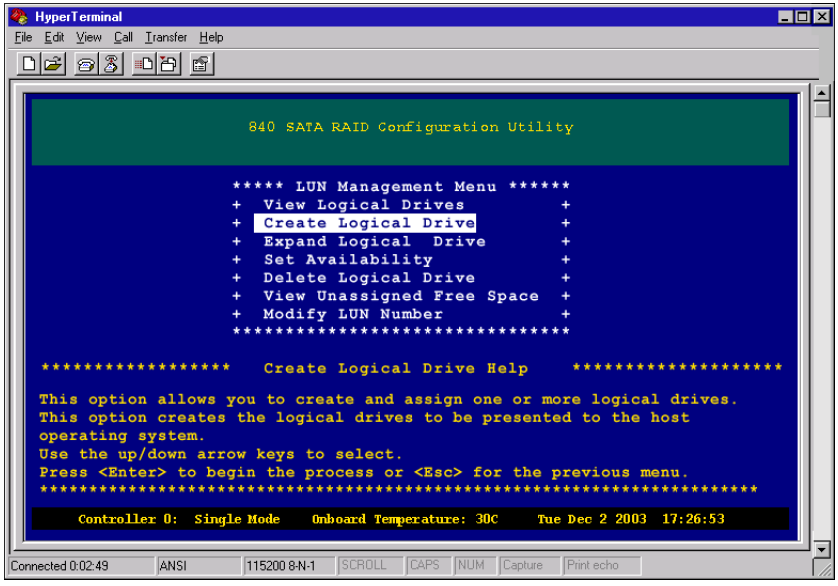

4 From the LUN Management menu, select **Create Logical Drive**, then press **ENTER**. The *Select Unused Region Menu* opens.

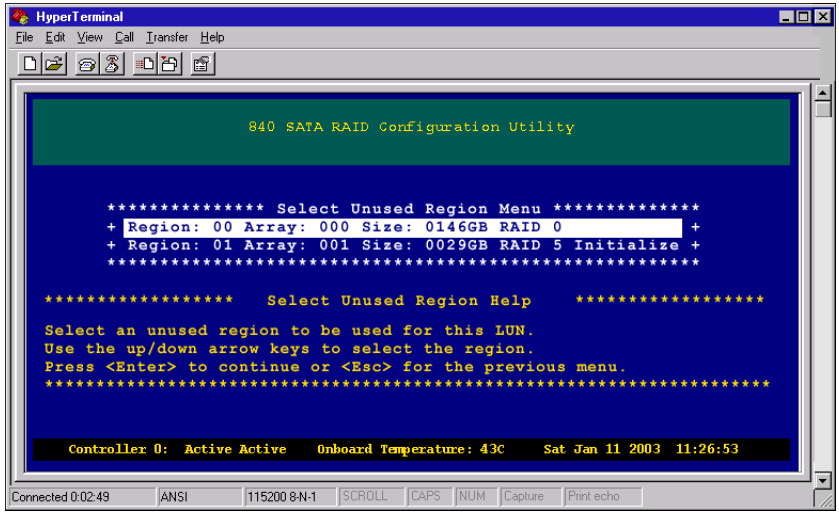

5 Press the up and down arrow keys to select a region you want to use for the logical drive, then press **ENTER**. The *Number of Logical Drives to create* screen opens.

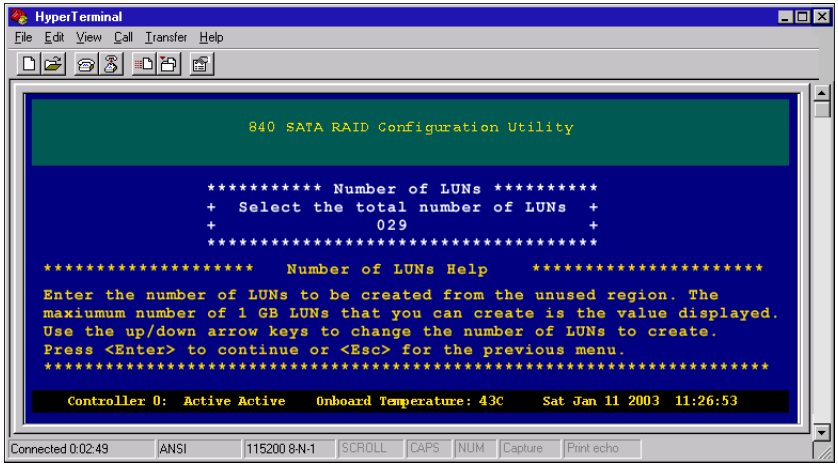

Important

The number shown as the default is the maximum number of 1 GB logical drives you can create using the selected free space region in the previous step. In this case the free space region was 29 GB which yielded a total of 29, 1 GB Host logical drives.

6 Press the up and down arrow keys to change the desired number of Host LUNs to create, then press ENTER. The *LUN Size* screen opens.

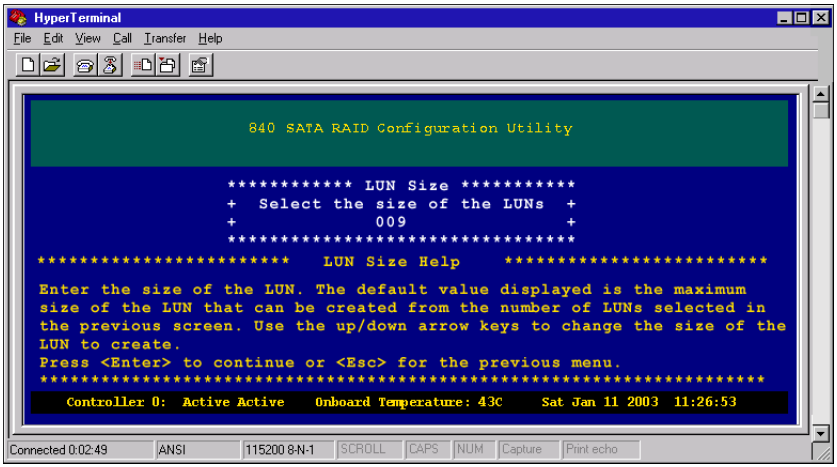

7 Press the up and down arrow keys to change the desired size of the LUN, then press ENTER.

The value displayed is the default size of the logical drive that can be created using the number of logical drives set in the previous screen based on the unused region space. If you are creating an even number of logical drives from an odd size value the remaining unused space is made available as unused space from which another separate logical drive can be created. For example, if you have a fault tolerant array with unused region of 143 GB and you create 2 logical drives from that unused space, you will have two logical drives each with 71 GBs. After you create the logical drives and view the unused region you will see the remaining 1 GB available for use in another logical drive.

8 When you are prompted to save the configuration, select **Yes**, then press ENTER.

9 Press any key to return to the Main Menu.

Expanding a logical drive

To expand a LUN:

1 From the Main Menu, select **Configuration Menu**, then press **ENTER**. The *Configuration Menu* opens.

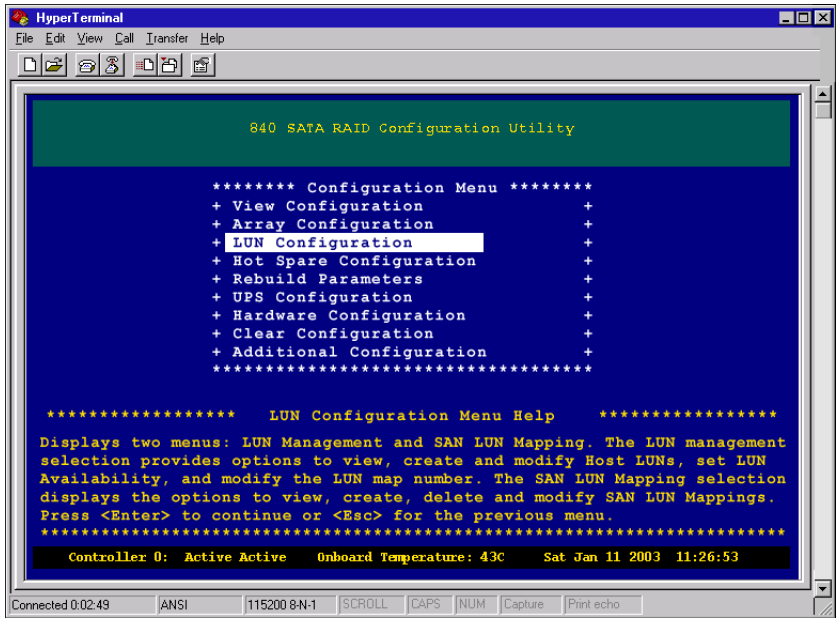

2 From the Configuration menu, select **LUN Configuration**, then press ENTER. The *LUN Configuration Menu* opens.

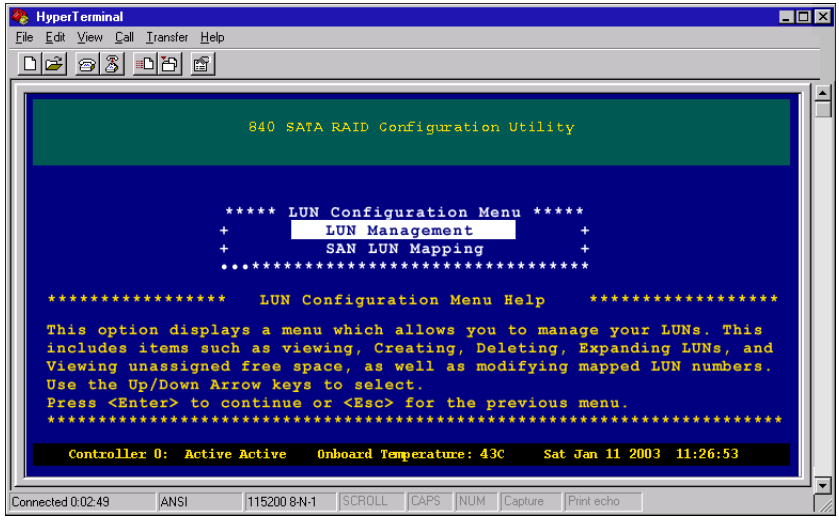

3 From the LUN Configuration Menu, select **LUN Management**, then press ENTER. The *LUN Management Menu* opens.

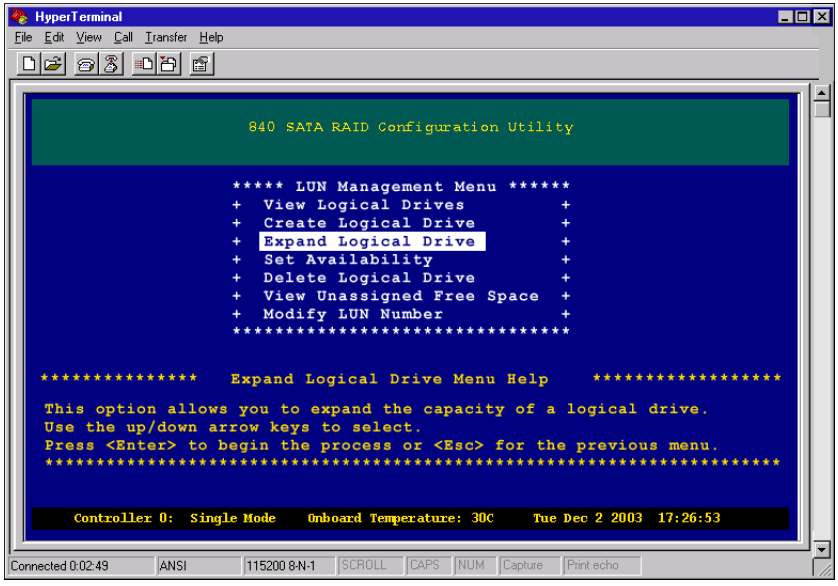

4 From the LUN Management menu, select **Expand Logical Drive**, then press **ENTER**. The *Select Logical Drive* screen opens.

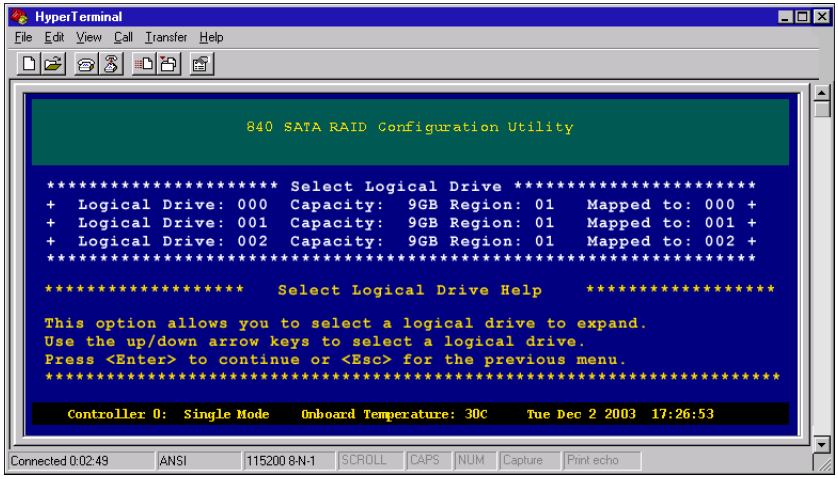

5 Select a logical drive that you want to expand, then press **ENTER**. The *Unused Regions For Expansion* screen opens.

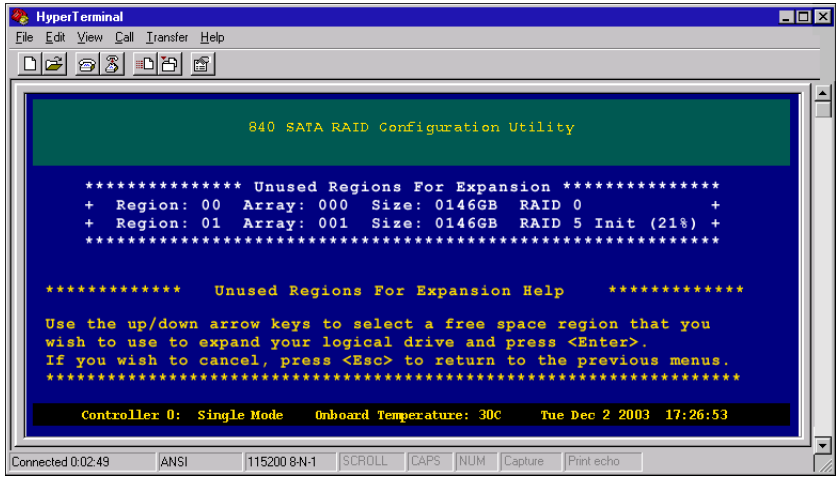

6 Select the unused region into which you want to expand the logical drive, then press ENTER. The *Expansion Size* screen opens.

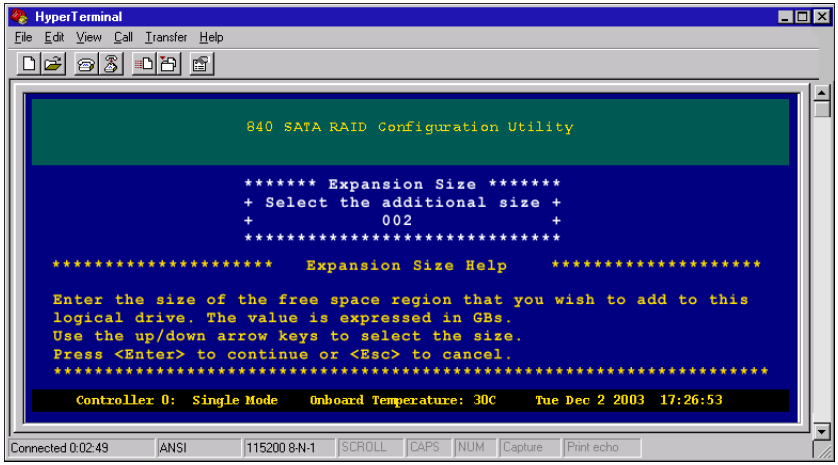

7 Press the up and down arrow keys to change the value to the desired size, then press ENTER.

8 When you are prompted to save the configuration, select **Yes**, then press ENTER.

9 Press any key to return to the Main Menu.

Setting logical drive availability

This option will make a logical drive available to a controller port and determines whether that logical drive is seen by all or a specific host system attached to that controller port.

To set logical drive availability:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

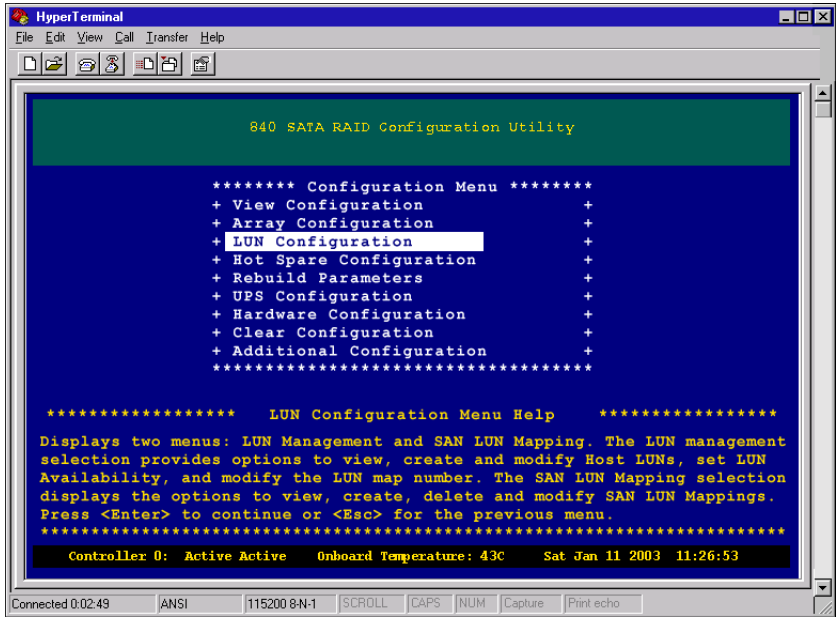

2 From the Configuration menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

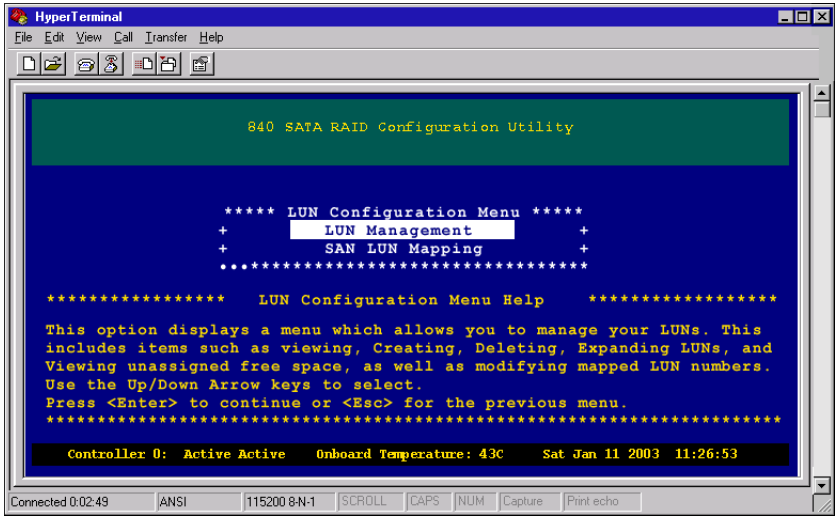

3 From the LUN Configuration menu, select **LUN Management**, then press **ENTER**. The *LUN Management Menu* opens.

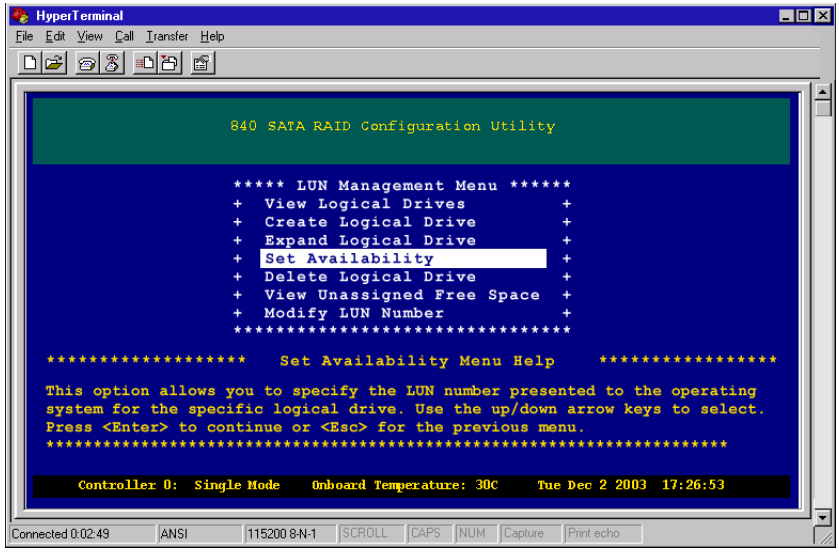

4 From the LUN Management Menu, select **Set Host LUN Availability**, then press **ENTER**. The *Select Logical Drive for Availability* screen opens.

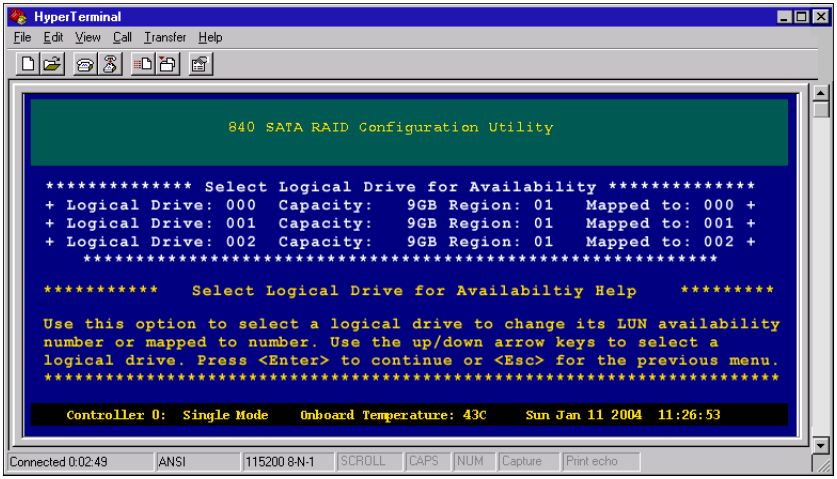

5 From the Select Logical Drive for Availability screen, select a logical drive that you want to map to another LUN number, then press **ENTER**. The *Modify the Logical Drive Availability* screen opens.

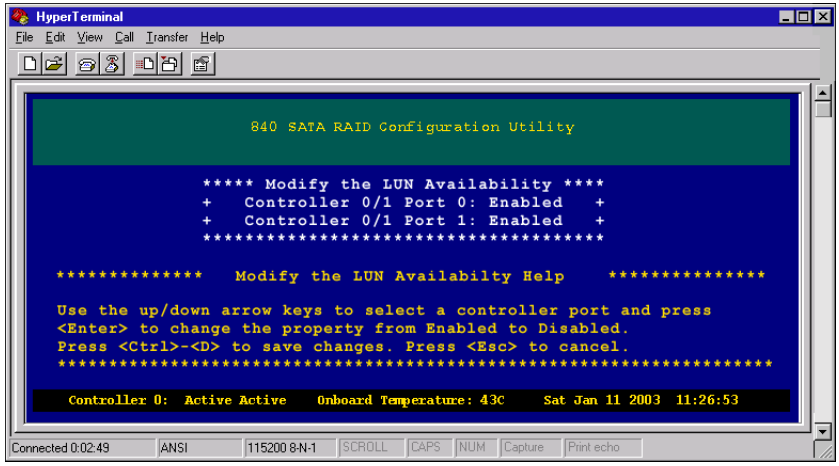

6 From the Modify the Logical Drive Availability screen, select the controller port to modify. Press **ENTER** to change the property from Enable to Disable.

7 If necessary, repeat the step to change the availability of the other port.

8 Press CTRL + D to save the changes.

9 When prompted to save the configuration, select **Yes**, then press ENTER.

10 Press any key to return to the Main Menu.

Deleteing a logical drive

This process will remove an existing logical drive.

To delete a logical drive:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

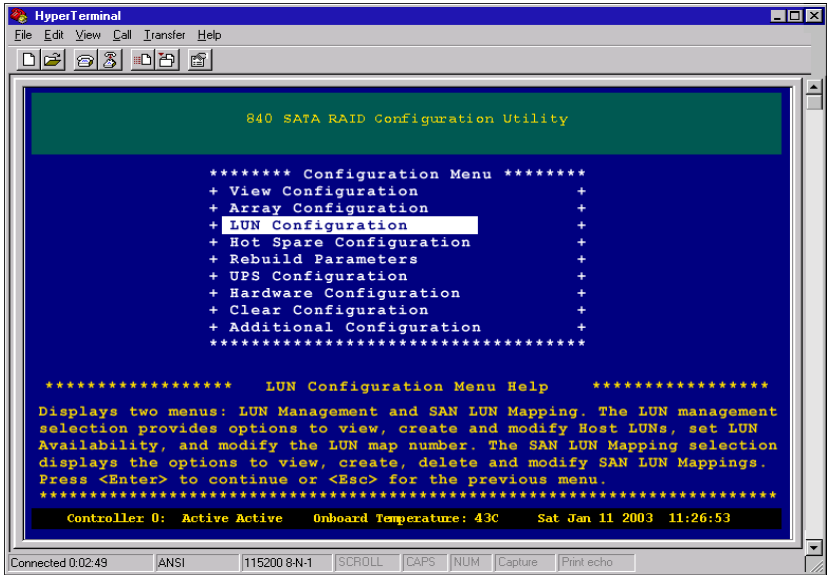

2 From the Configuration menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

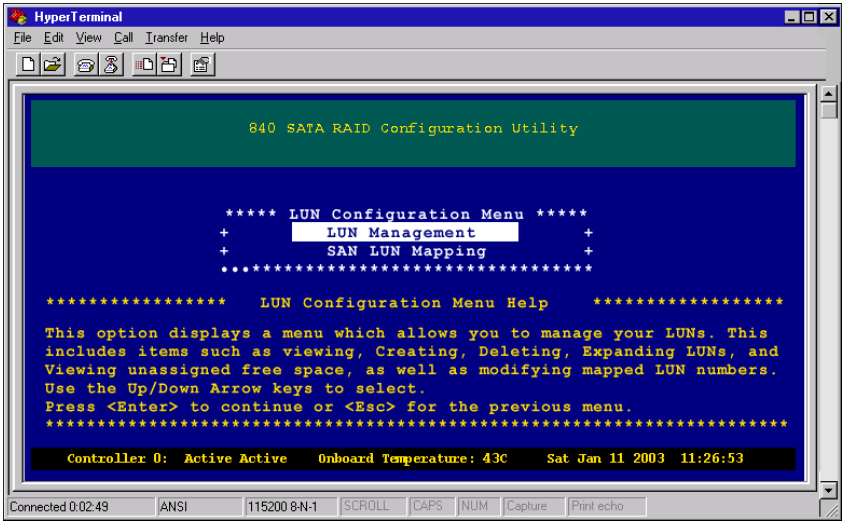

3 From the LUN Configuration menu, select **LUN Management**, then press **ENTER**. The *LUN Management Menu* opens.

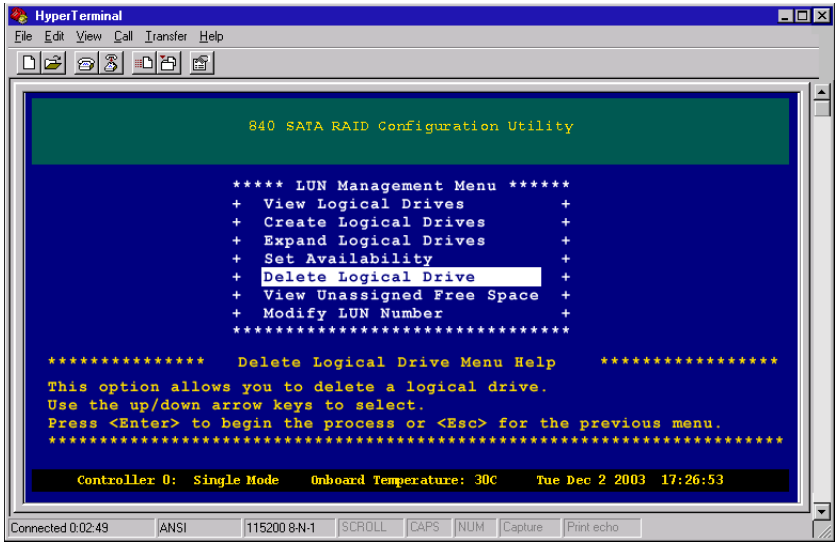

4 From the LUN Management Menu, select **Delete Logical Drive**, then press **ENTER**. The *Select a Logical Drive to Delete* screen opens.

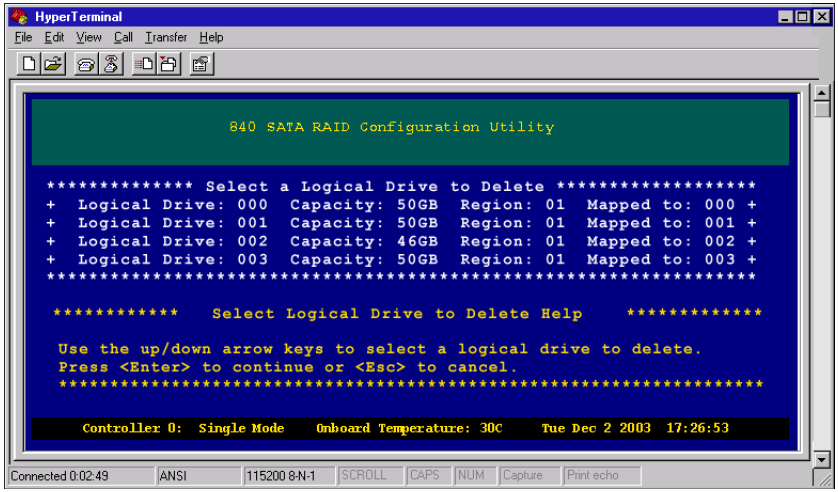

5 From the Select a Logical Drive to Delete screen, select the logical drive to delete, then press **ENTER**.

6 When you are prompted to save the configuration, select **Yes**, then press **ENTER**.

7 Press any key to return to the Main Menu.

Modifying a mapped LUN

This option lets you change the assigned LUN number after the logical drive has previously been made available.

To modify a LUN number:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

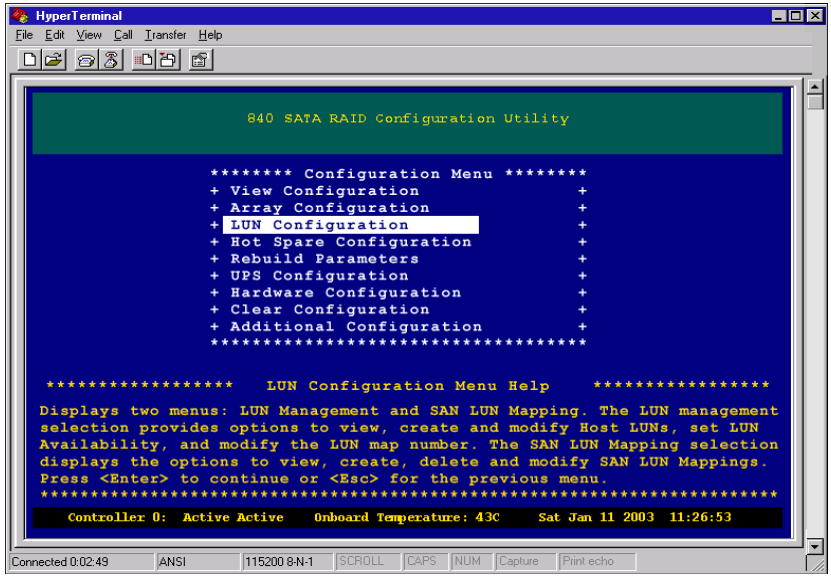

2 From the Configuration menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

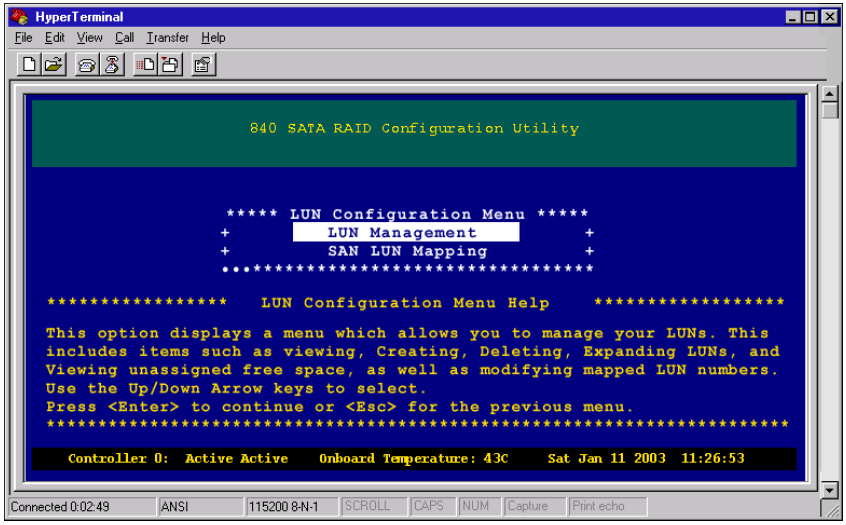

3 From the LUN Configuration menu, select **LUN Management**, then press **ENTER**. The *LUN Management Menu* opens.

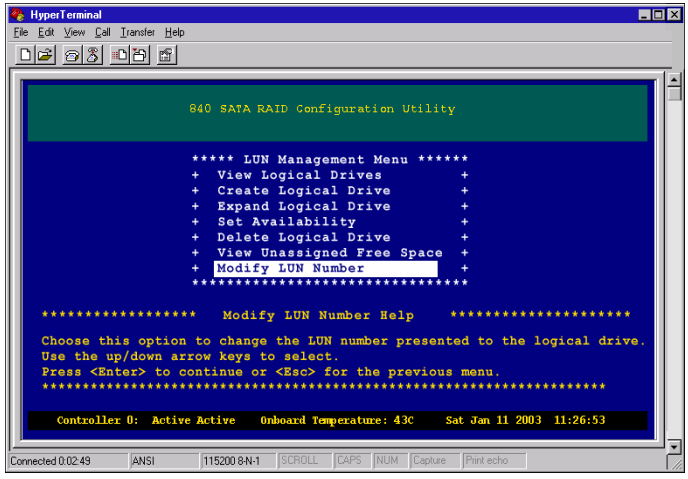

4 From the LUN Management menu, select **Modify Mapped LUN Number**, then press **ENTER**. The *Select a Logical Drive to Modify Mapped Number* screen opens.

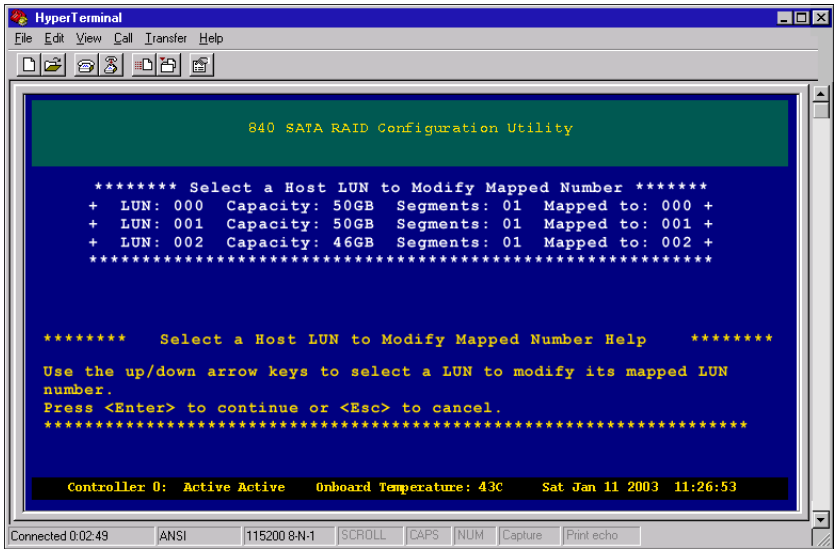

5 From the Select a Logical Drive to Modify Mapped Number screen, select the logical drive whose number you intend to change, then press **ENTER**. The *Select a New LUN Number* screen opens.

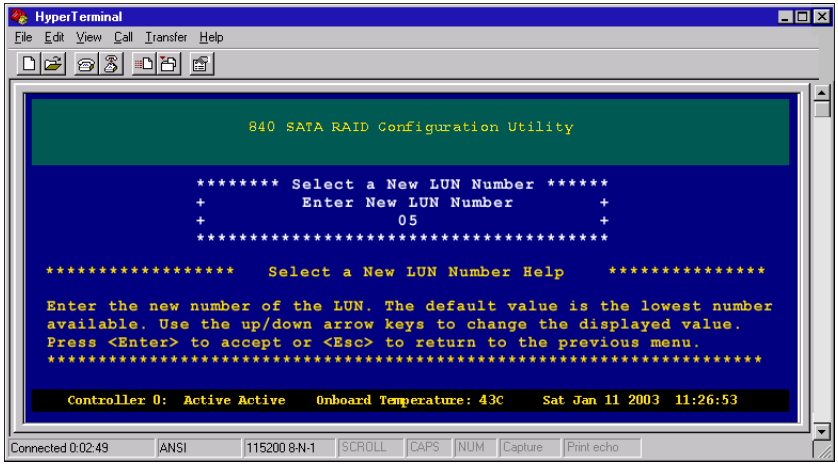

6 Press the up and down arrow keys to change the logical drive number to the desired number, then press ENTER.

7 When you are prompted to save the configuration, select **Yes**, then press ENTER.

8 Press any key to return to the Main Menu.

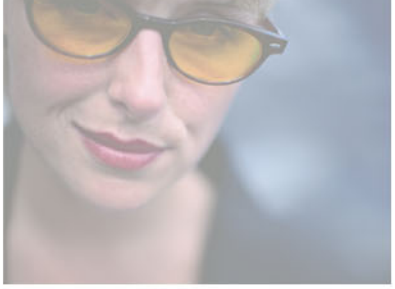

SAN LUN Mapping

5

This chapter provides information on SAN LUN mapping. Read this chapter to learn how to:

- Access the SAN LUN mapping configuration menu
- View SAN LUN mappings
- Create a SAN LUN mapping
- Delete a SAN LUN mapping
- Modify a SAN LUN mapping
- View connected hosts

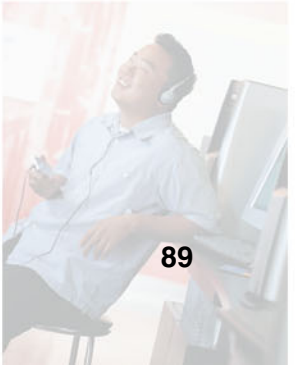

Overview

When attaching multiple host systems in a SAN environment, it may be necessary to more precisely control which hosts have access to which logical drives. In addition to controlling logical drive availability on a RAID controller on a port-by-port basis, it is also possible to further restrict access to a specific logical drive. Up to 512 SAN LUN mappings are supported.

Terminology

The following table describes the terminology relating to SAN LUN Mapping.

Term	Description
Node Name	This is an eight-byte hexadecimal number, uniquely identifying a single host system. It incorporates the World Wide Name and two additional bytes which are used to specify the format.
Port Name	This is an eight-byte hexadecimal number, uniquely identifying a single host port. It incorporates the World Wide Name and two additional bytes which are used to specify the format and indicate the port number.
Mapping Name	A 32-character name that can be used to help identify the host system.
Exclusive Access	A logical drive is presented to only one host system. It is not available to any other host systems.
Read/Write Access	A logical drive will allow both reads and write operations.
Read Only Access	A logical drive will not allow writes.
Mapped LUN Number	This is the LUN number that a specific logical drive responds to when accessed by a host. It is not necessary for this to bear any relation to the LUN number.
Mapping Availability	Which controller ports the mapping is valid for.

Accessing the SAN LUN mapping configuration menu

To access the SAN LUN mapping configuration menu:

1 From the Main Menu, select the **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

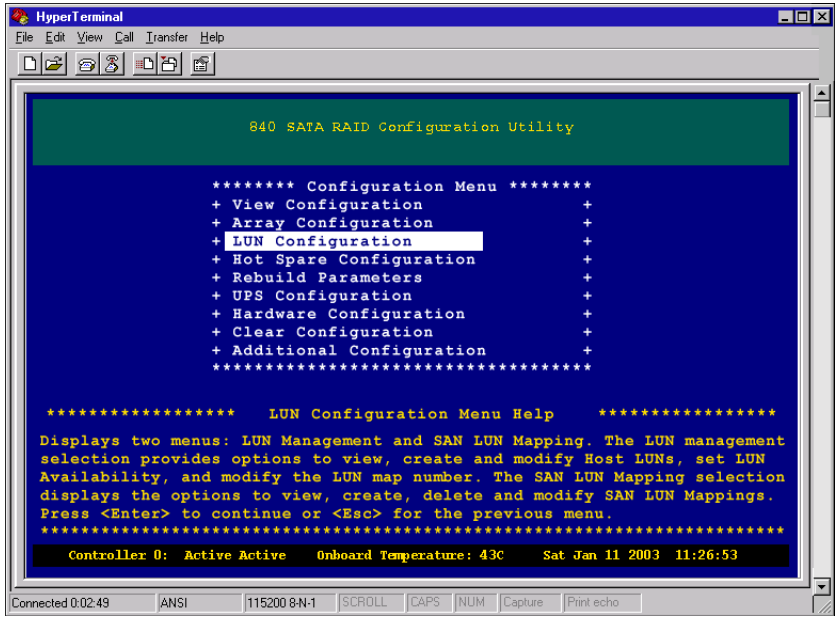

2 From the Configuration Menu, select **LUN Configuration**, then press ENTER. The *LUN Configuration Menu* opens.

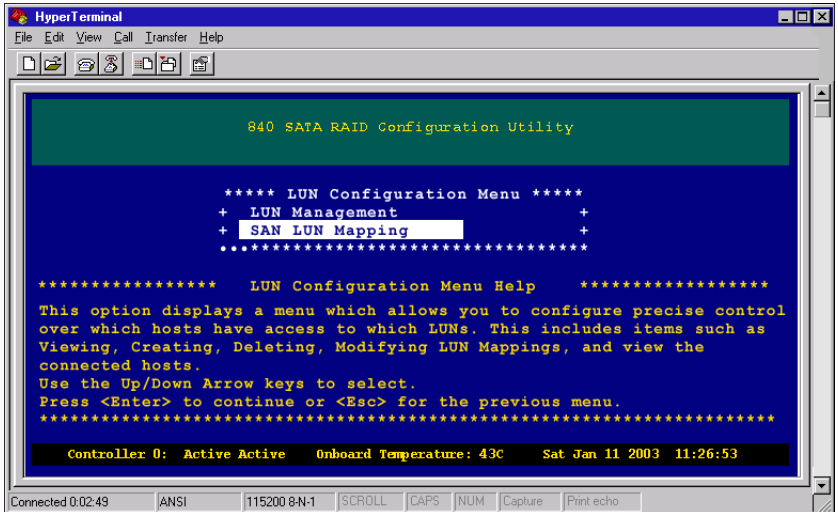

3 From the LUN Configuration menu, select **SAN LUN Mapping**, then press ENTER. The *SAN LUN Mapping* screen opens.

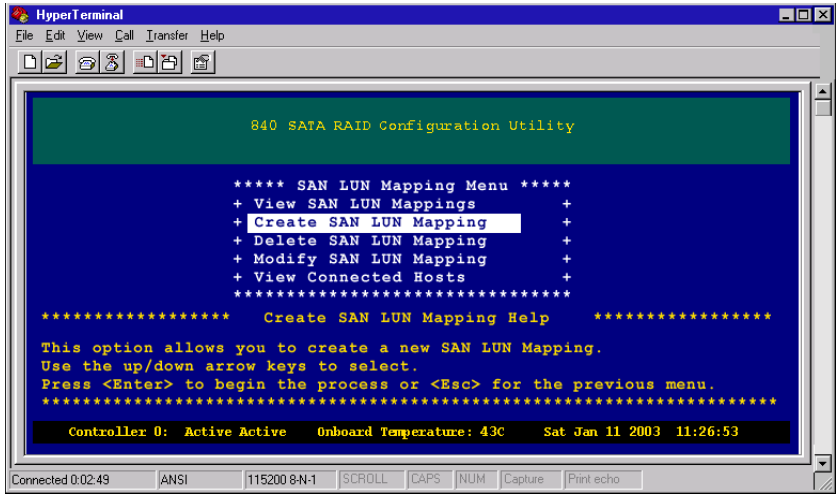

The SAN LUN Mapping Configuration Menu provides the options for viewing, creating, deleting, and managing SAN LUNs, and for viewing the host systems attached to the storage network.

Viewing SAN LUN mappings

To view existing SAN LUN mappings, choose this option.

To view SAN LUN mappings:

1 From the Main Menu, select the **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

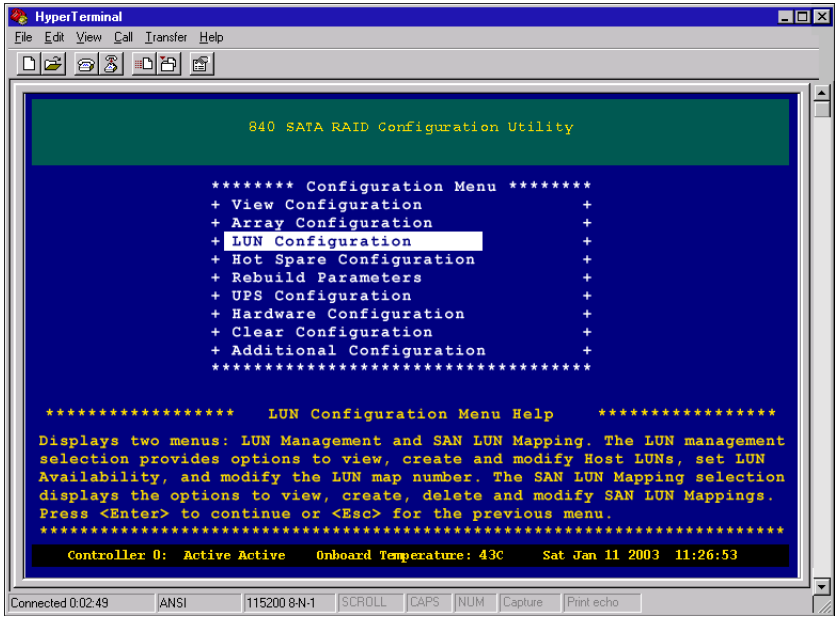

2 From the Configuration menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

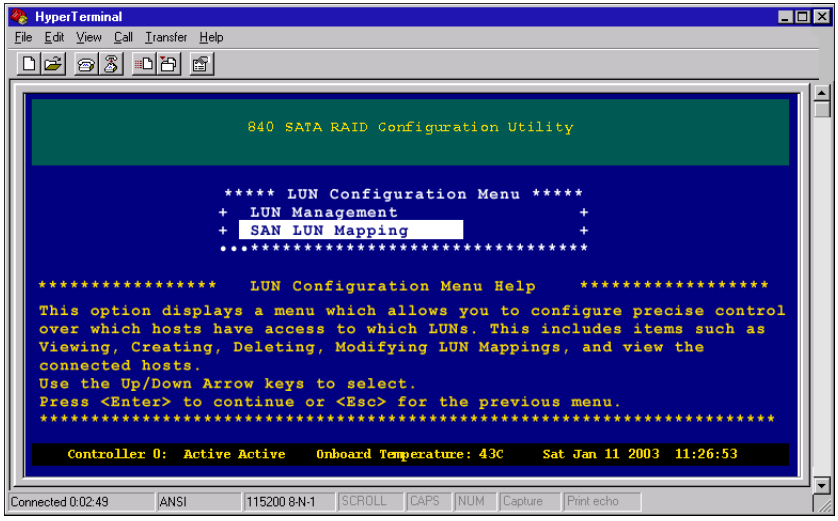

3 From the LUN Configuration menu, select **SAN LUN Mapping**, then press **ENTER**. The *SAN LUN Mapping* screen opens.

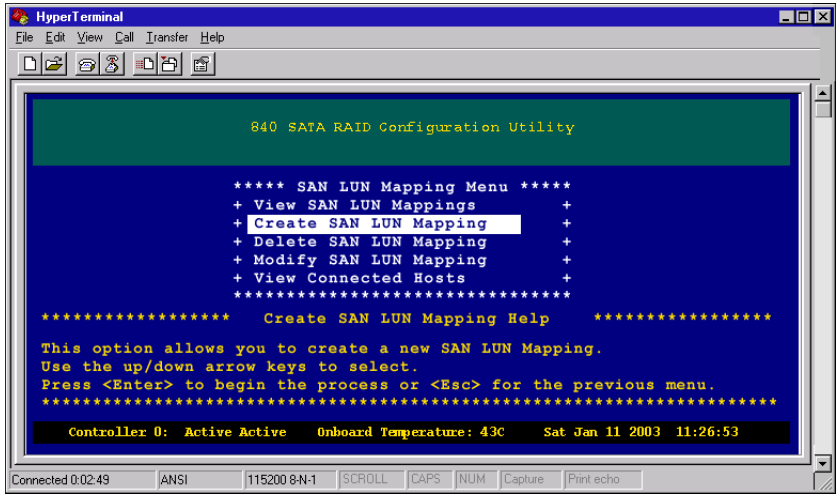

4 From the SAN LUN Mapping Menu, select **View SAN LUN Mappings**, then *p*ress **ENTER**. The *View SAN LUN Mapping* screen opens.

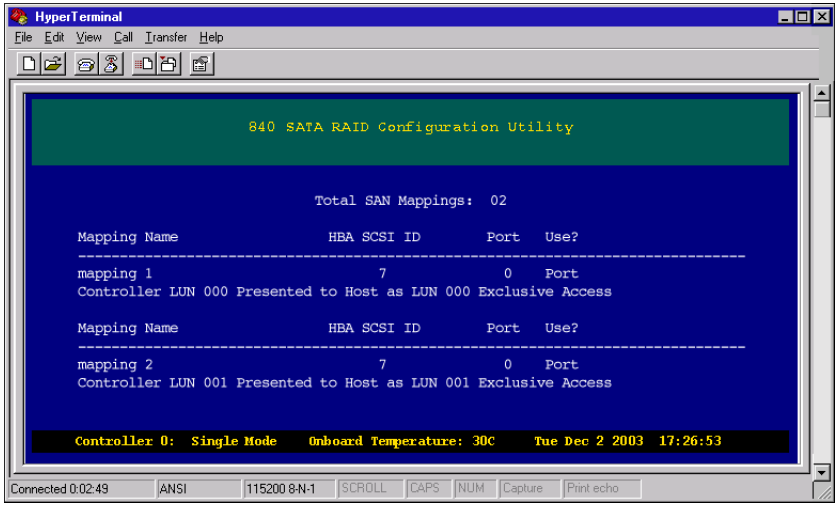

In the example above, there are two mappings created. They have an identical node name but different port names, indicating they are both in the same host system. The port name is used for the mapping, and access is only allowed on Port 0 of the controllers. Each Host HBA port can access one logical drive, exclusively.

The following information is shown on the screen.

- **Mapping Name:** This is the name defined by the user when creating the SAN LUN Mapping. You may use up to 32 ASCII characters.
- **HBA SCSI ID:** This is initiator (target) SCSI ID of the host system HBA port used for the mapping.
- **Port:** This is the port of the RAID Controller. The controller has two ports, and will be listed here as “0” and “1.” They represent “P0” and “P1” respectively.
- **Description:** Below each mapping is a plain language description of how the Controller LUN is being presented to the host system as, and what access rights are established for the mapping. These will be Exclusive, Shared, Read-Only.

Creating a SAN LUN mapping

To create a SAN LUN mapping:

1 From the Main Menu, select the **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

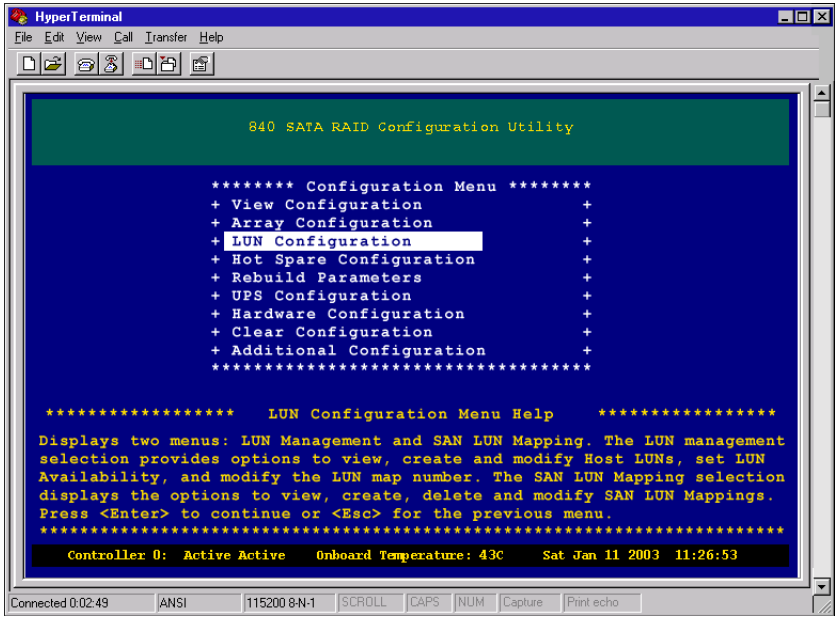

2 From the Configuration menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

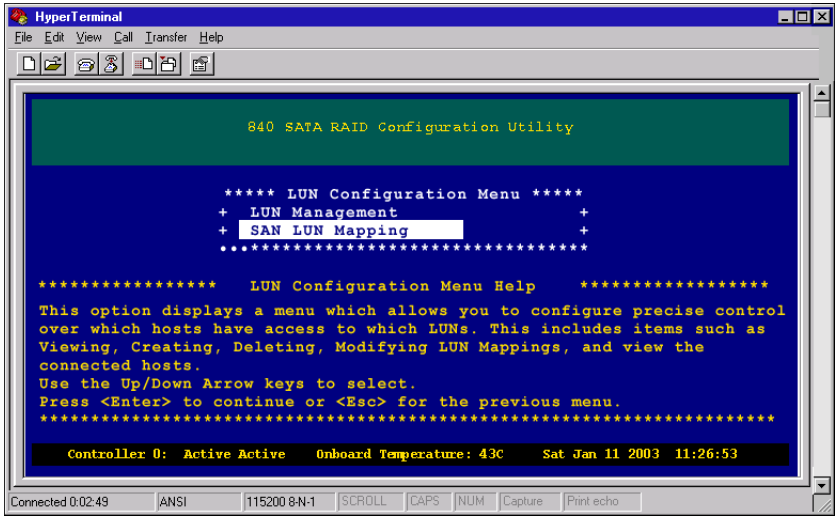

3 From the LUN Configuration menu, select **SAN LUN Mapping**, then press **ENTER**. The *SAN LUN Mapping* screen opens.

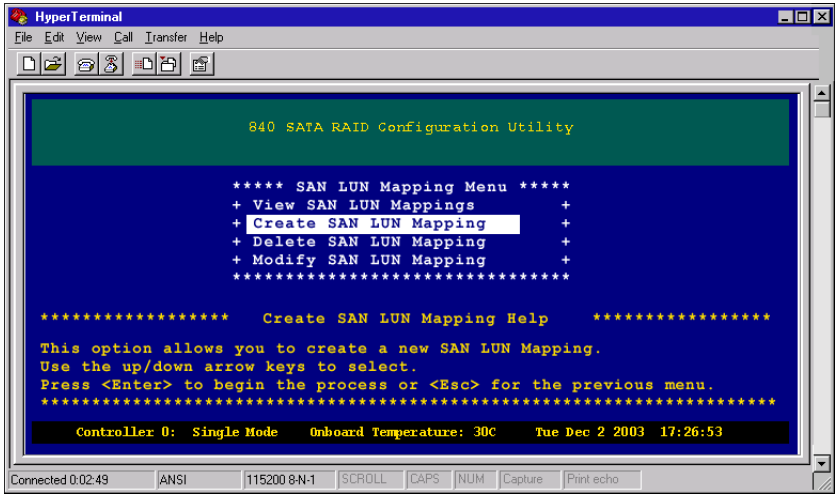

4 From the SAN LUN Mapping Menu, select **Create SAN LUN Mapping**, then press **ENTER**. The *SAN LUN Mapping Name* screen opens.

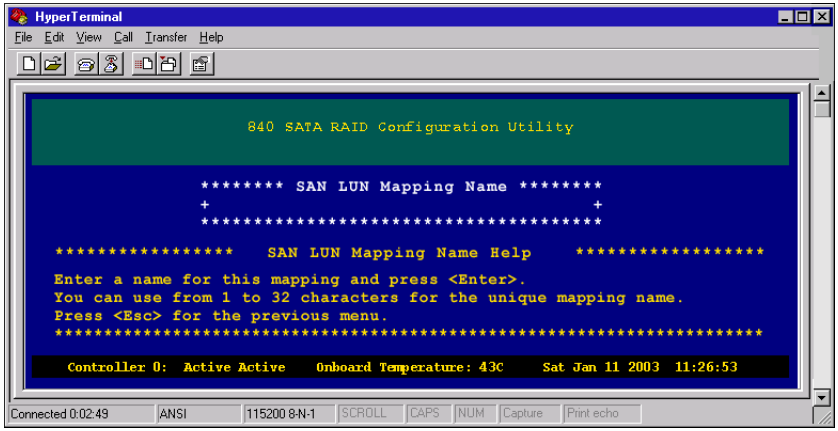

5 Type a unique name for this mapping (as many as 32 characters), then press **ENTER**. The *SAN LUN Mapping Configuration* screen will open.

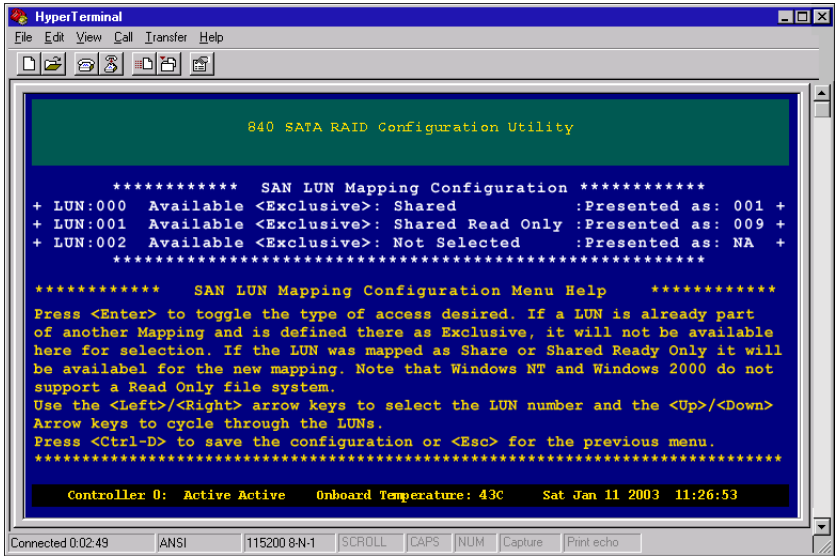

If a LUN is already part of another mapping and is defined as Exclusive, it will not be available here for selection. If the logical drive was mapped as Share or Shared Read Only, it will be available.

Important

Microsoft Windows NT, Microsoft Windows 2000, and Microsoft Windows 2003 do not support a Read Only file system.

6 When you are prompted to save the configuration, select **Yes**, then press **ENTER**.

7 Press any key to return to the Main Menu.

Deleting a SAN LUN mapping

To delete a SAN LUN mapping:

1 From the Main Menu, select the **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

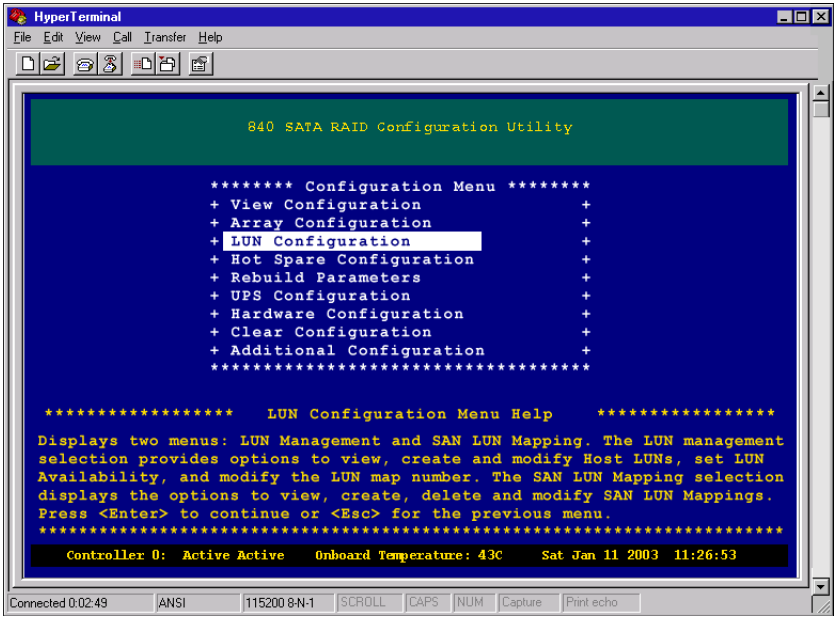

2 From the Configuration Menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

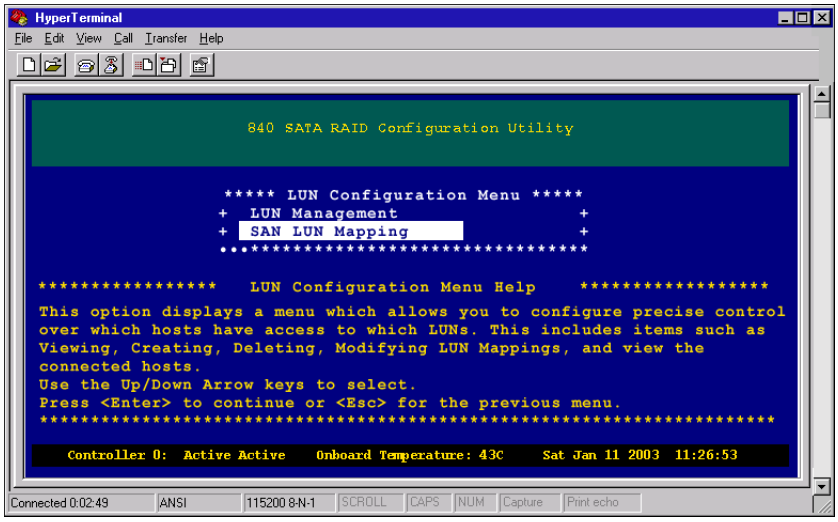

3 From the LUN Configuration Menu, select **SAN LUN Mapping**, then press **ENTER**. The *SAN LUN Mapping* screen opens.

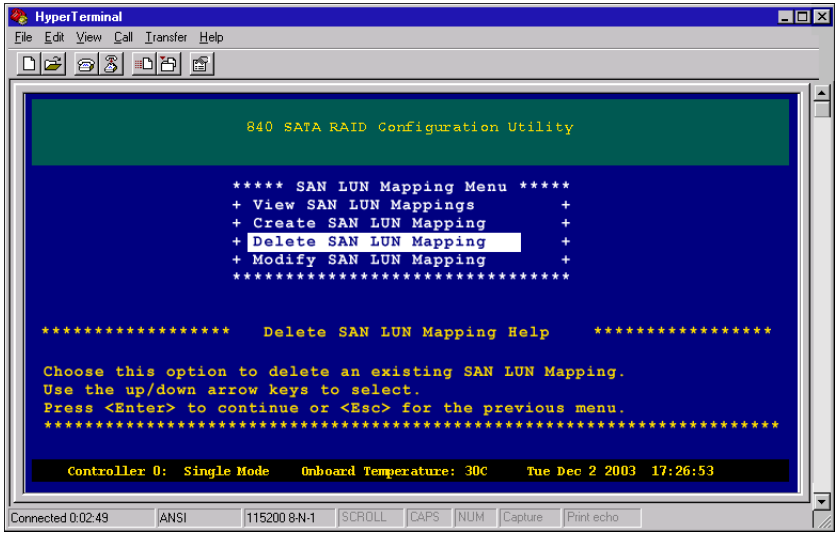

4 From the SAN LUN Mapping menu, select **Delete SAN LUN Mapping**, then press **ENTER**. The *Delete SAN LUN Mapping* screen opens.

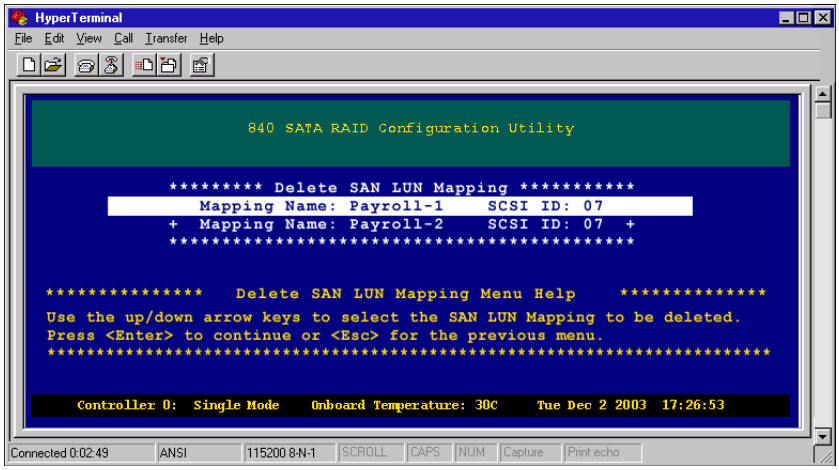

5 Select a SAN LUN Mapping to be deleted, then press **ENTER**.

Tips & Tricks

SAN LUN Mappings can be deleted in any order without affecting the other mappings.

6 When you are prompted to save the configuration, select **Yes**, then press **ENTER**.

7 Press any key to return to the Main Menu.

8 Access your host operating system and make the necessary adjustments for the change in disk(s).

Modifying SAN LUN Mapping

It is sometimes necessary to modify some or all of the parameters of a SAN LUN Mapping. You can change the Mapping Name, Port Name, and the RAID controller port availability access of the mapping.

This may be necessary if, for example, a host system has an adapter failure and a new adapter with different Node and Port names (initiator ID for SCSI HBAs) is installed.

To modify a SAN LUN mapping:

1 From the Main Menu, select the **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

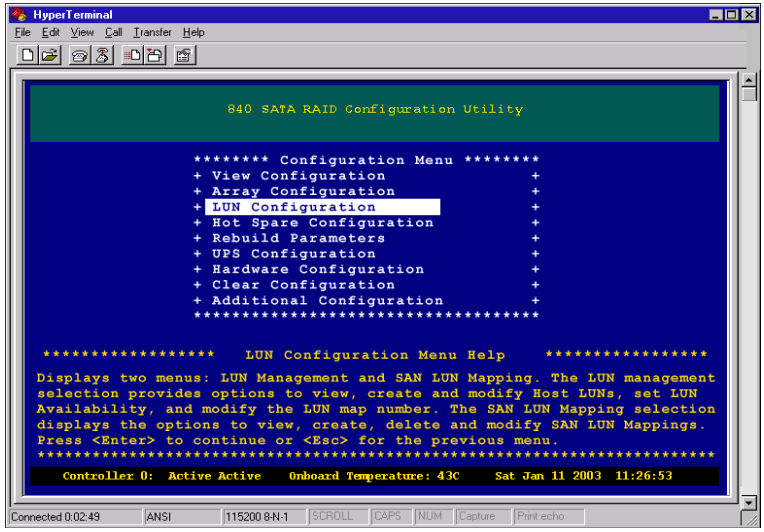

2 From the Configuration menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

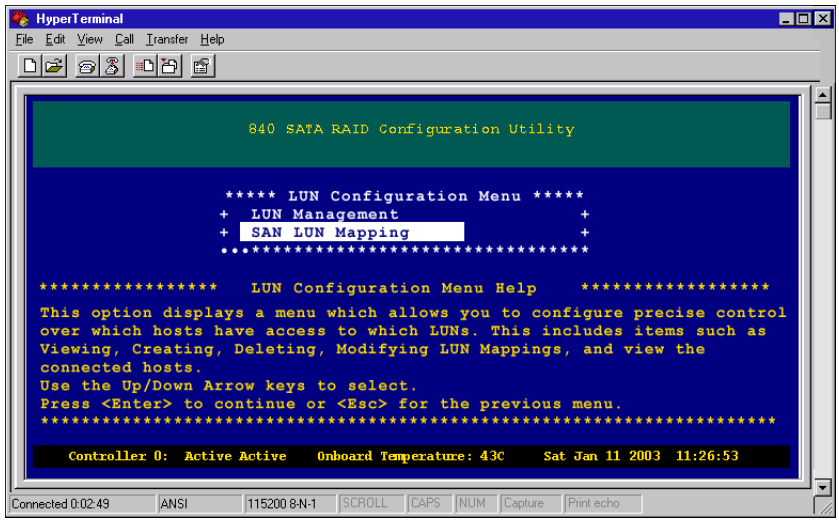

3 From the LUN Configuration menu, select **SAN LUN Mapping**, then press **ENTER**. The *SAN LUN Mapping* screen opens.

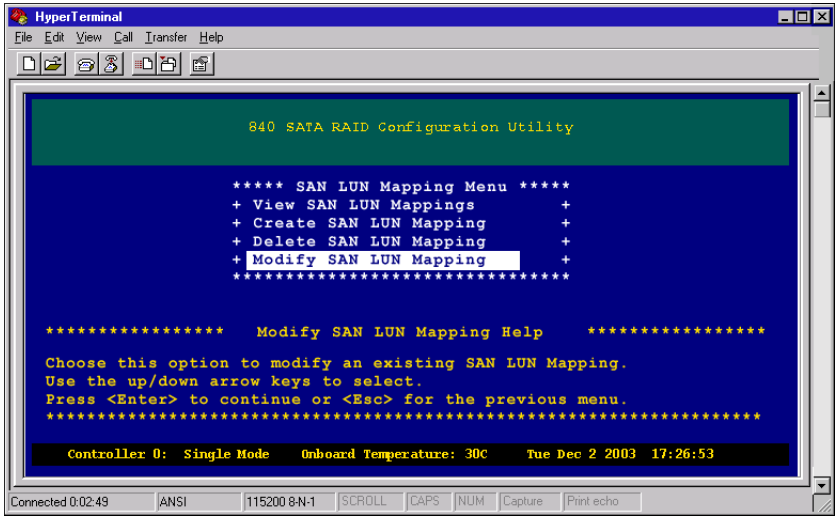

4 From the SAN LUN Mapping menu, select **Modify SAN LUN Mapping**, then press **ENTER**. The *Modify SAN LUN Mapping* screen opens.

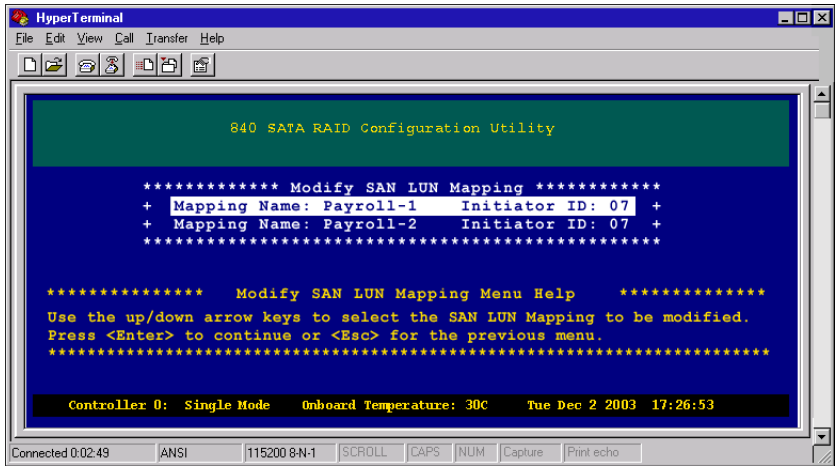

5 Press the up and down arrow keys to select a SAN LUN Mapping to be modified, then press **ENTER**. The *Modify SAN LUN Mapping Name* screen opens.

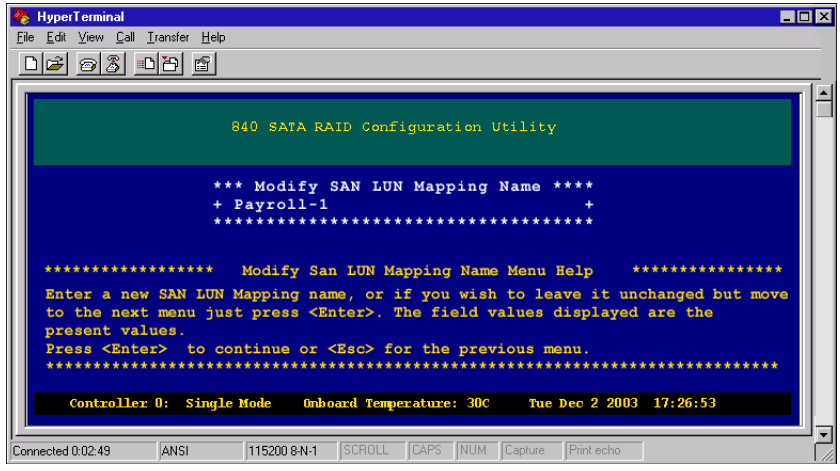

6 Enter a new SAN LUN mapping name, then press **ENTER**. The *Modify SAN LUN Mapping Configuration* screen opens.

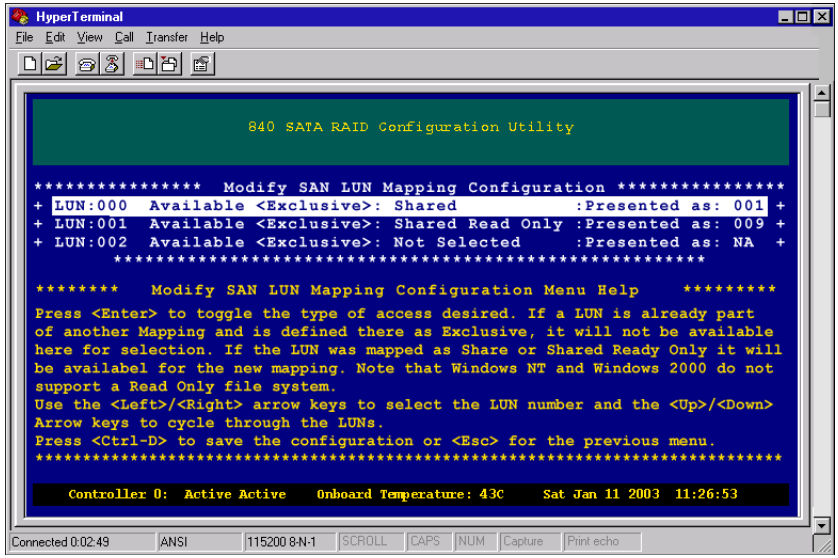

7 Change the parameters for the SAN LUN Mapping. Set the parameters for the SAN LUN Mapping by using the **ENTER** key to toggle the type of access desired (shared, shared read only, exclusive, or NA). Press the left or right arrow keys to change the LUN number and the up or down arrow keys to cycle through the logical drives.

If a logical drive is already part of another mapping, and is defined as Exclusive, it will not be available here for selection. If the logical drive was mapped as Share or Shared Read Only, it will be available.

Press **CTRL** + **D** to move to the next menu.

8 You are presented with modifying the controller LUN for the mapping. You can choose to leave the property in the automatic mode, or set an exclusive Controller LUN number for this mapping by choosing **User Defined** and entering a new LUN number.

Press **ENTER** to change the mode. If you selected **User Defined**, use the up or down arrow keys to change the Controller LUN number.

Press **CTRL** + **D** to move to the next menu.

9 When you are prompted to save the configuration, select **Yes**, then press ENTER.

10 Press any key to return to the Main Menu.

Viewing connected hosts

Choose this option to view all connected host HBA ports.

To view connected hosts:

1 From the Main Menu, select the **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

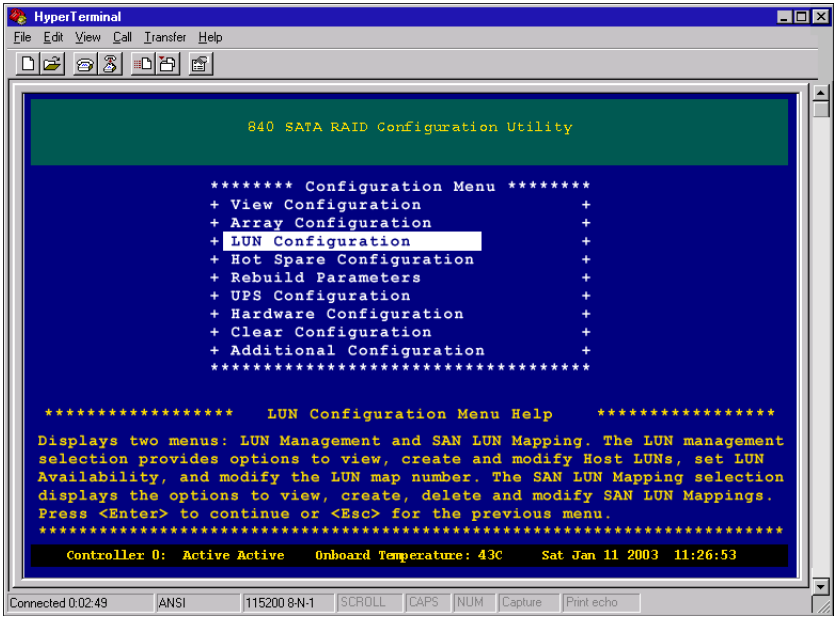

2 From the Configuration menu, select **LUN Configuration**, then press **ENTER**. The *LUN Configuration Menu* opens.

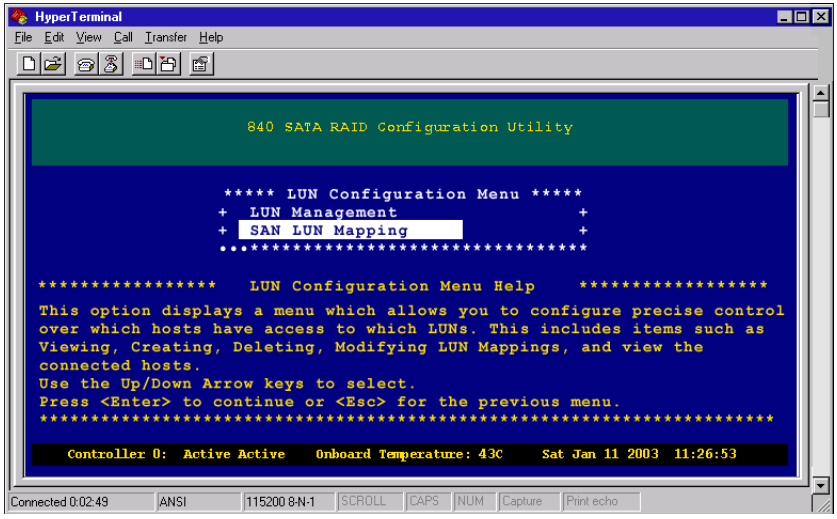

3 From the LUN Configuration menu, select **SAN LUN Mapping**, then press **ENTER**. The *SAN LUN Mapping* screen opens.

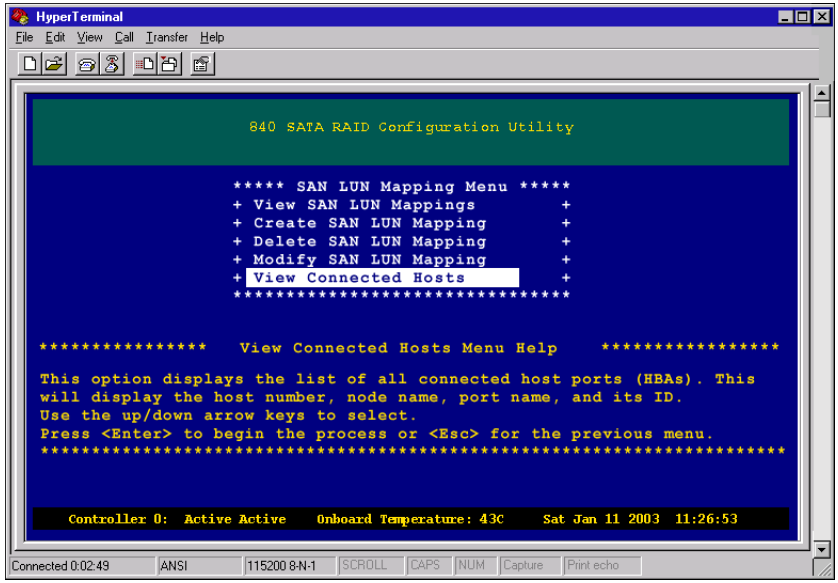

4 From the SAN LUN Mapping Menu, select **View Connected Host**, then press **ENTER**. The *View Logged-in Hosts* screen opens.

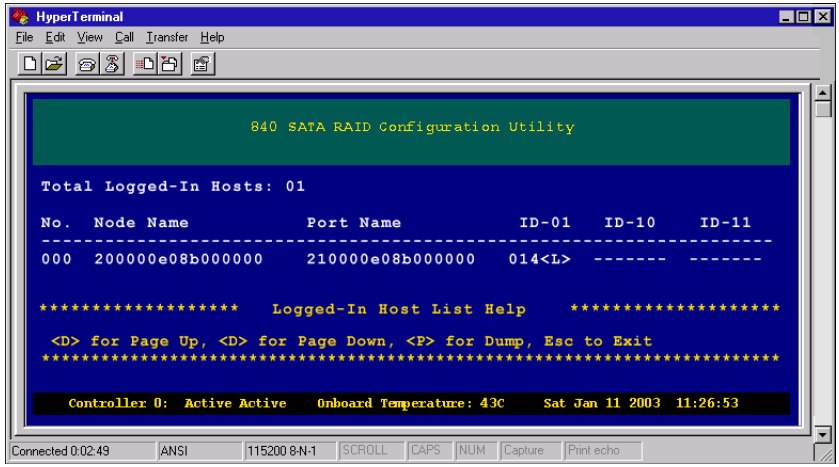

The following table provides an explanation of each of the fields shown in the display.

Field	Description
No.	Host number assigned by the controller. This is used when creating a SAN LUN Mapping and selecting a host from the list.
SCSI ID	This is the SCSI ID of the host bus adapter.

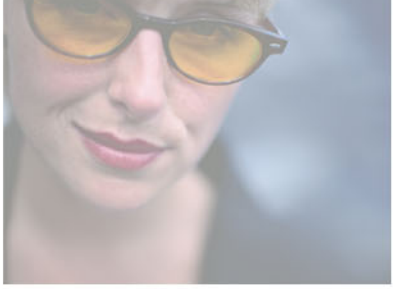

6 Advanced Configuration and Management

This chapter provides information on advanced configuration and management. Read this chapter to learn how to:

- Set rebuild parameters
- Configure hot spares
- Configure hardware settings
- View controller information
- Clear the configuration

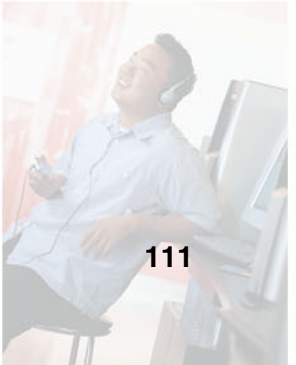

Setting rebuild parameters

The Rebuild Parameters menu option lets you enable or disable automatic rebuild operations, set the rebuild priority, and configure the controller for the auto hot spare feature.

Enabling/disabling auto rebuild

Whenever a hot spare is created, whether it is a global or dedicated spare (for the specific array), the automatic rebuild feature is enabled. However, there may be occasions where it is necessary to manually enable or disable this option.

When enabled, auto regeneration will begin the rebuild cycle when a fault tolerant array is configured, a hot spare drive is available, and a drive fault occurs.

To enable or disable auto rebuild:

1 From the Main Menu select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

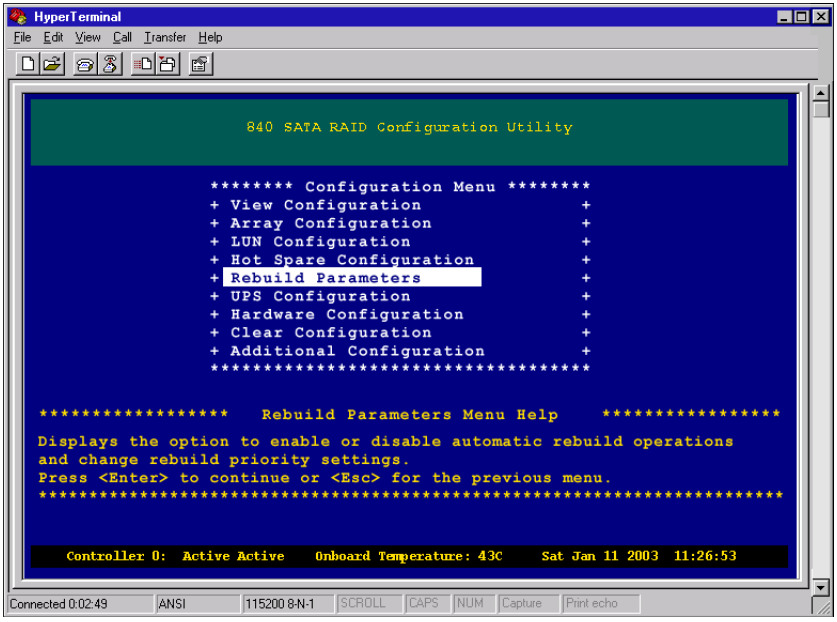

2 From the Configuration menu, select **Rebuild Parameters**, then press ENTER. The *Rebuild Parameters* menu opens.

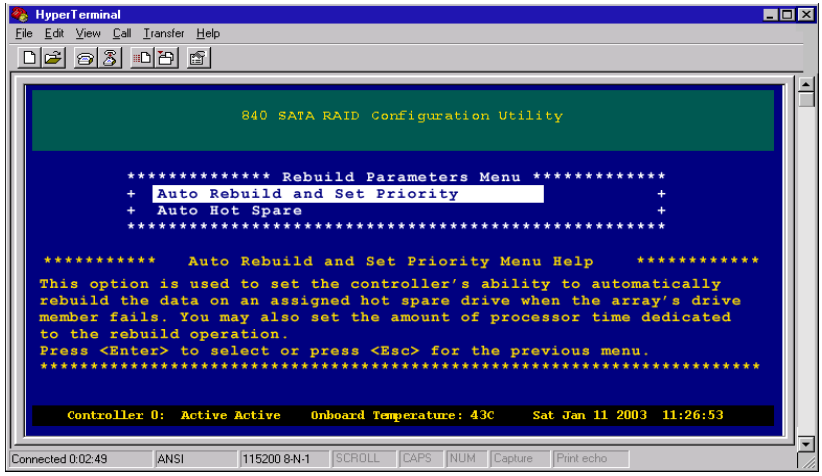

3 From the Rebuild Parameters menu, select the **Auto Rebuild and Set Priority** option, then press ENTER. The *Auto Rebuild and Set Priority* menu opens.

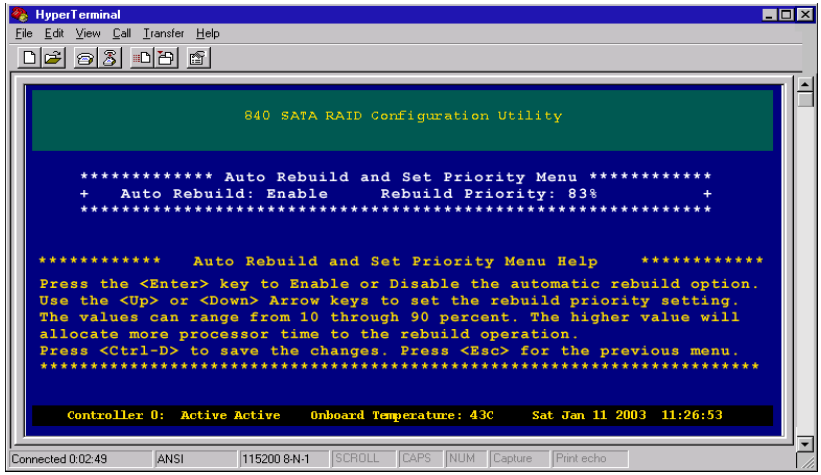

From the Auto Rebuild and Set Priority menu you have the option to enable the automatic rebuild, or set the percent of processor time allocated to a rebuild operation.

4 Press ENTER to **Enable** or **Disable** Auto Rebuild. Then press the up and down arrow keys to change the rebuild priority value.

Important

The values range from 10% through 90%. The higher the value the more processor time is allocated to the rebuild operation.

5 Press CTRL + D to save the changes and continue.

6 When you are prompted to confirm the save, choose **Yes**, then press ENTER.

7 Press any key to return to the Main Menu.

Configuring auto hot spare

This feature lets you configure the controller to automatically start a rebuild operation when a replacement drive is inserted into the drive slot of the removed failed drive member. This feature only applies when a dedicated or pool spare has not been assigned.

To configure an auto hot spare:

1 From the Main Menu select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

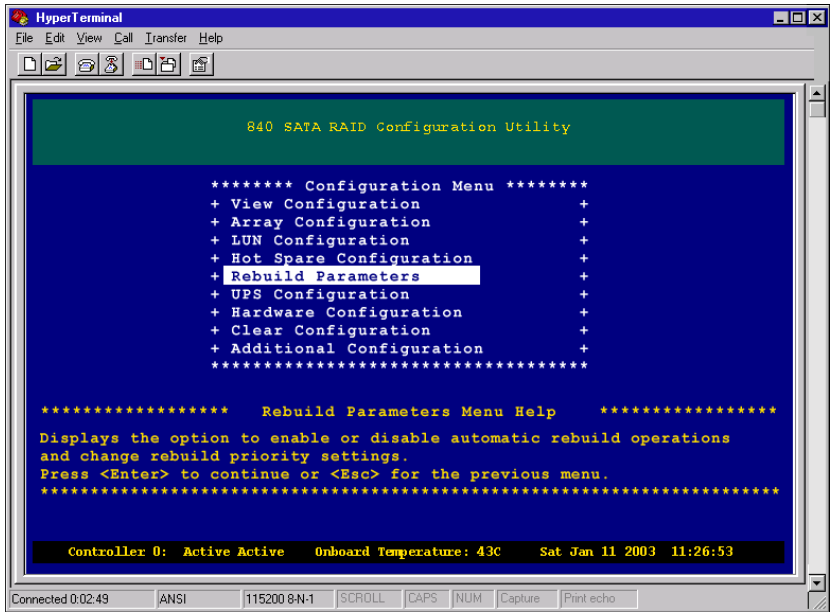

2 From the Configuration Menu, select **Rebuild Parameters**, then press **ENTER**. The *Rebuild Parameters Menu* opens.

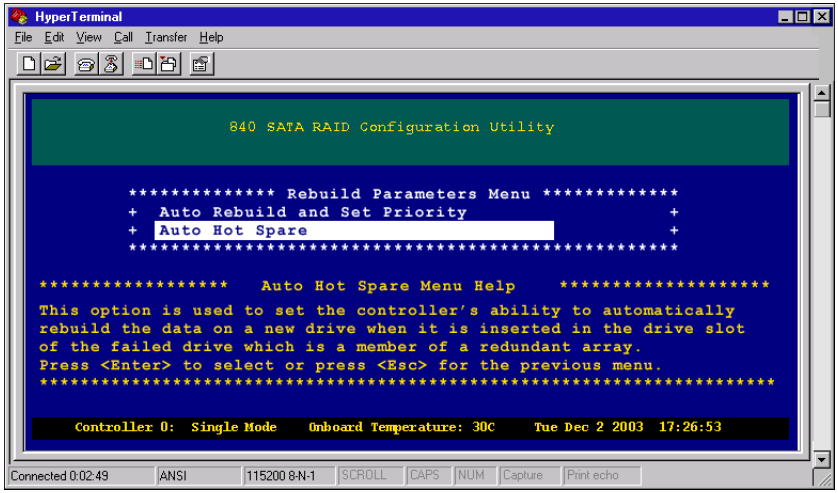

3 From the Rebuild Parameters Menu, select the **Auto Hot Spare** option, then press **ENTER**. The *Auto Hot Spare Menu* opens.

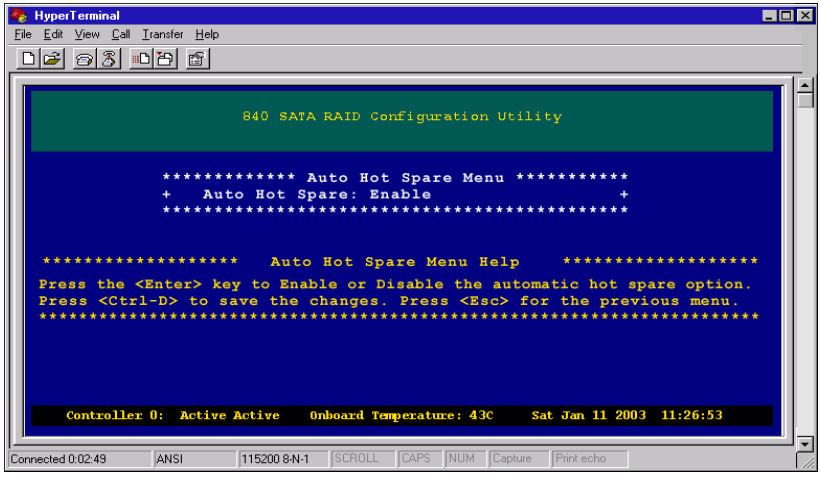

4 Press ENTER to toggle between **Enable** or **Disable**. Enabling this feature will automatically start a rebuild operation when a replacement drive is inserted into the drive slot of the removed failed drive member.

5 Press CTRL + D to save the changes or press ESC to cancel and return to the previous menu.

Configuring hot spares

The process of configuring fault tolerant arrays includes assigning drives as global or dedicated hot spares. In the event of a drive failure that is a member of a fault tolerant array, the controller will use an assigned global hot spare to replace the failed drive in any array. If a dedicated spare is assigned to the specific array, that array will use its dedicated hot spare to replace the failed drive member.

Viewing the list of hot spare drives

To view a list of hot spare drives:

1. From the Main Menu select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

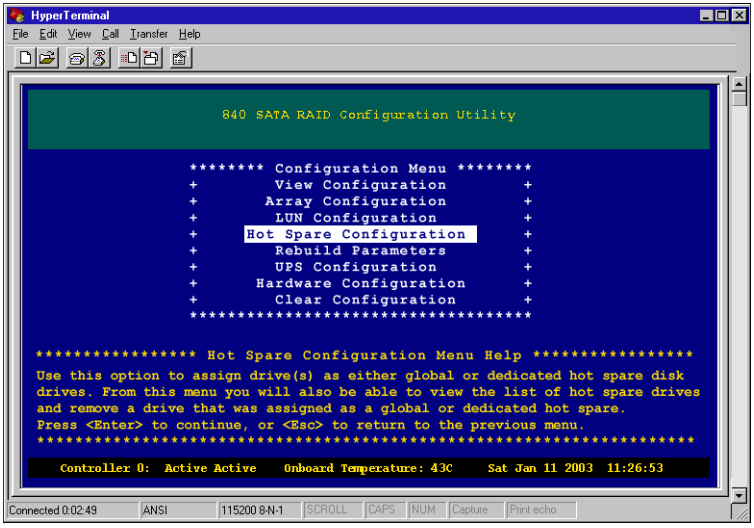

2 From the Configuration menu, select **Hot Spare Configuration**, then press **ENTER**. The *Hot Spare Configuration* menu opens.

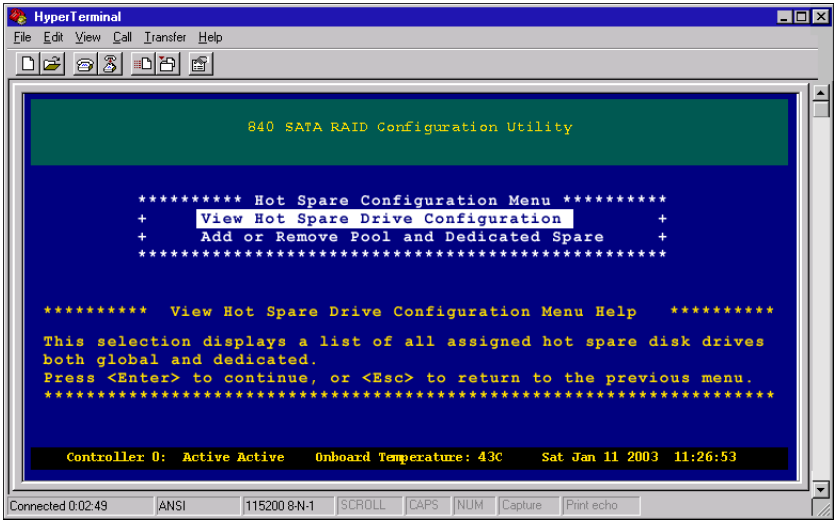

3 Select **View Hot Spare Drive Configuration,** then press ENTER. A list of hot spare drives opens.

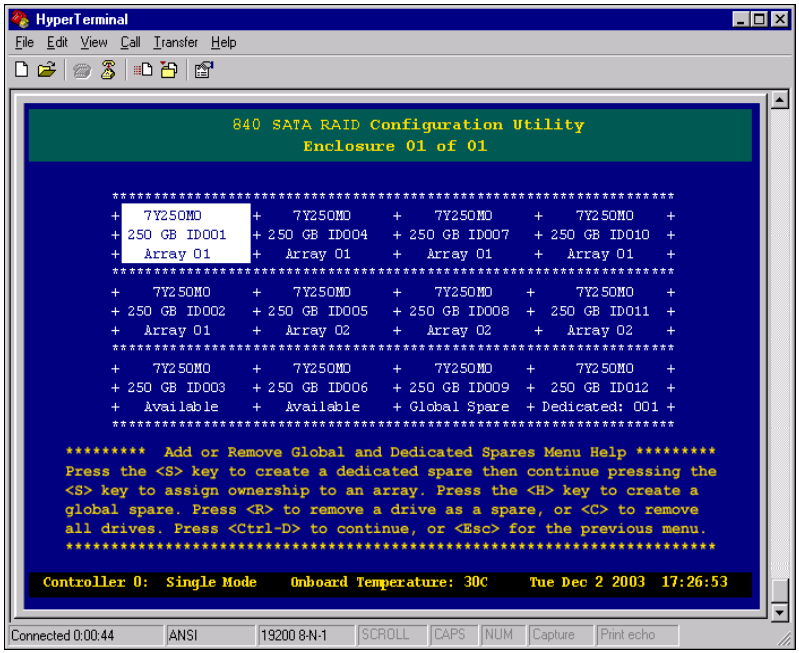

Adding or removing hot spare drives

To add or remove hot spare drives:

1 From the Main Menu select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

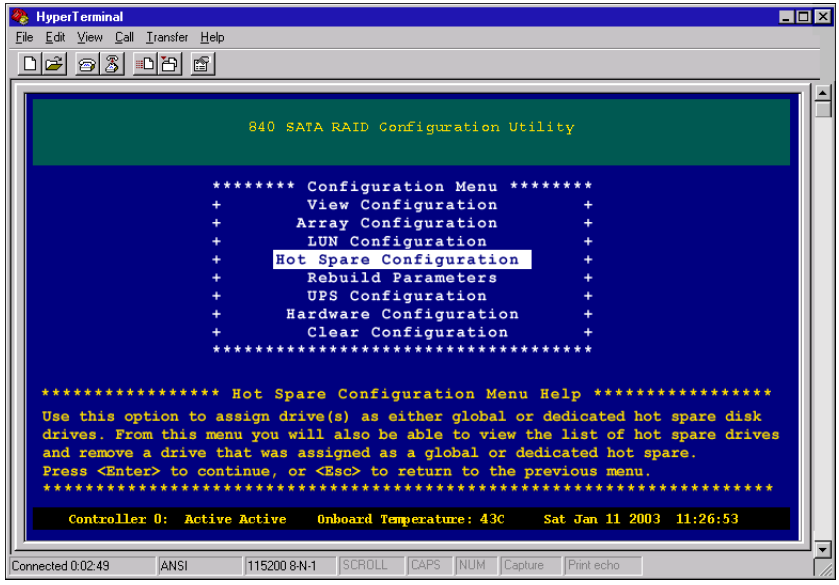

2 From the Configuration menu, select **Hot Spare Configuration**, then press ENTER. The *Hot Spare Configuration* menu opens.

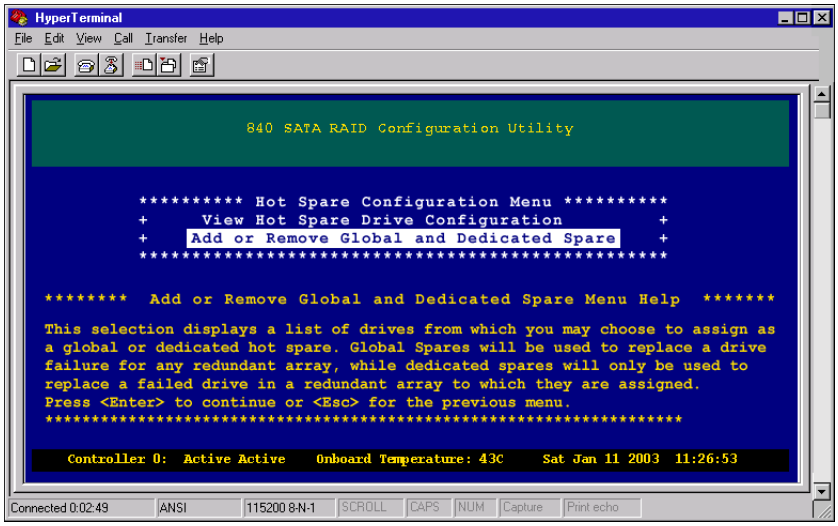

3 Select **Add or Remove Global and Dedicated Spare,** then press ENTER. The *Add or Remove Global and Dedicated Spares* screen opens.

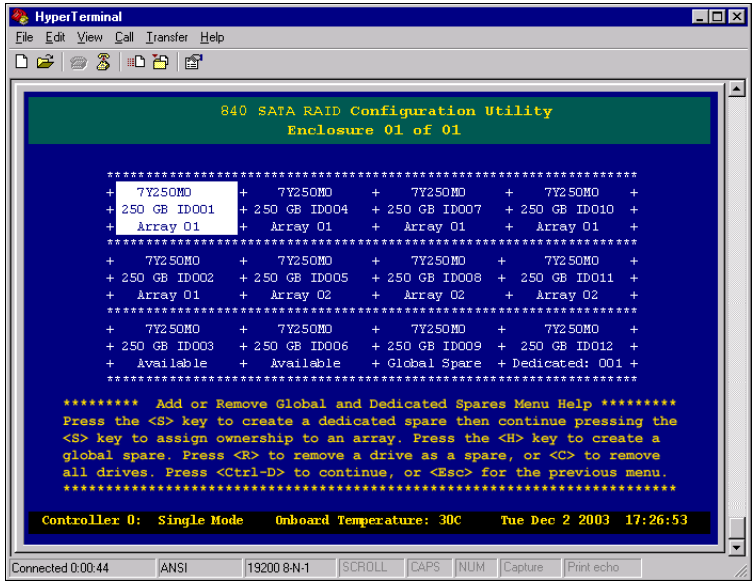

4 Press the arrow keys to select an **available** drive to use:

   To add a dedicated spare, press **S** to assign that drive as a dedicated spare. Then continue to press **S**, which displays the arrays from which you can assign the dedicated spare.

   To add a drive as a global spare, press the arrow keys to select an **available** drive, then press **H** key.

   To remove a drive as a hot spare, select the drive, then press **R**, or press **C** to remove all drives in the enclosure that are currently assigned as spares. After the changes have been made, press the **CTRL** + **D** to continue.

5 When you are prompted to save the changes, press **Y** to save the new hot spare configurations.

6 Press any key to continue.

Configuring the hardware

The hardware configuration menu provides the options to set the Controller LUN, configure the controller port connection options, set the controller port data rate, configure the controller's RS-232 port BAUD rate, and set the Packetized Data Transfer/QAS option.

Setting the controller LUN

This option lets the you override the automatic feature which causes the Controller LUN to automatically use the next available data LUN to communicate between StorView and the SATA RAID Controller. We recommend that this setting remain in the automatic mode.

Changing this parameter can alleviate a potential issue that some operating systems have with a controller showing up at the end of the data LUNs.

In the event you have an operating system that is having a problem with the controller LUN being displayed, select **Disabled**. This will allow the controller to use the next available data LUN but not be displayed to the operating system.

If the Controller LUN is required to be at a specific, user-defined LUN, choose the **User Defined** option and select from 0-511.

To configure the controller LUN:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

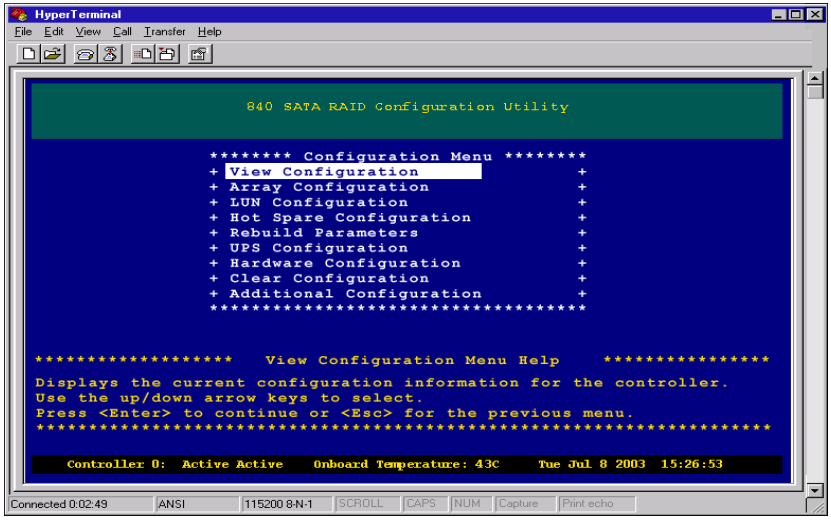

2 From the Configuration Menu, select **Hardware Configuration**, then press ENTER. The *Hardware Configuration Menu* opens.

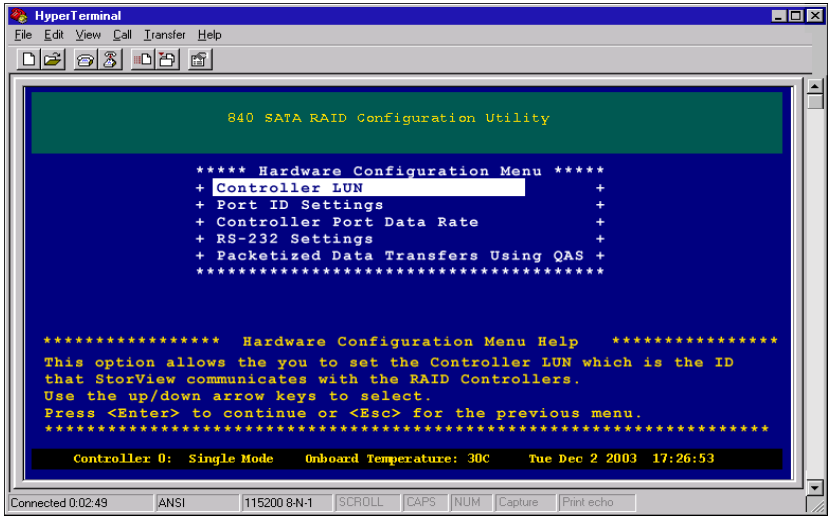

3 From the Hardware Configuration Menu, select **Controller LUN**, then press **ENTER**. The *Set Controller LUN* screen opens.

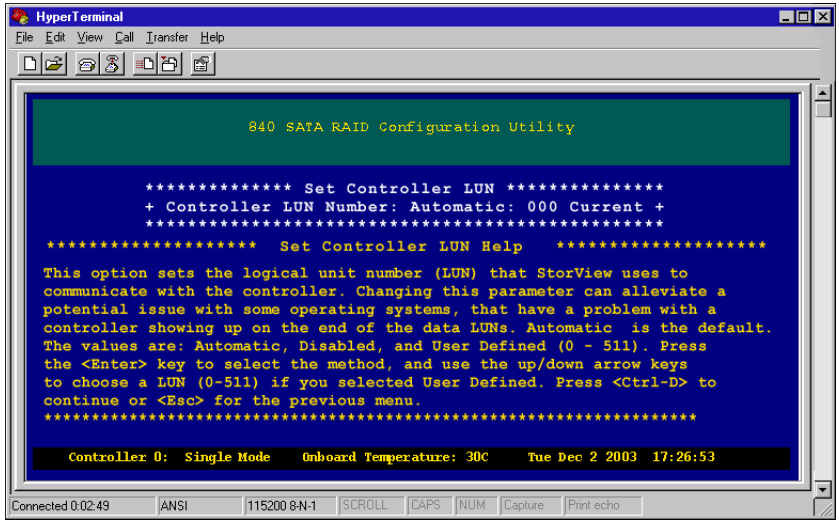

4 Press **ENTER** to select the method of Controller LUN assignment.

Important

If you chose **User Defined**, a specific LUN value is displayed and selected by the up or down arrow keys.

If you select a LUN number that is used by another device, the word “Conflict” will appear next to the LUN number.

5 Press **CTRL** + **D** to save the changes and continue.

6 When you are prompted to confirm the save, choose **Yes**, then press **ENTER**.

7 Press any key to return to the Main Menu.

Setting the controller port ID

The Port ID settings provide you the option to set the SCSI ID for the controller port. You can use IDs 0-15.

To set the ID for controller ports:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

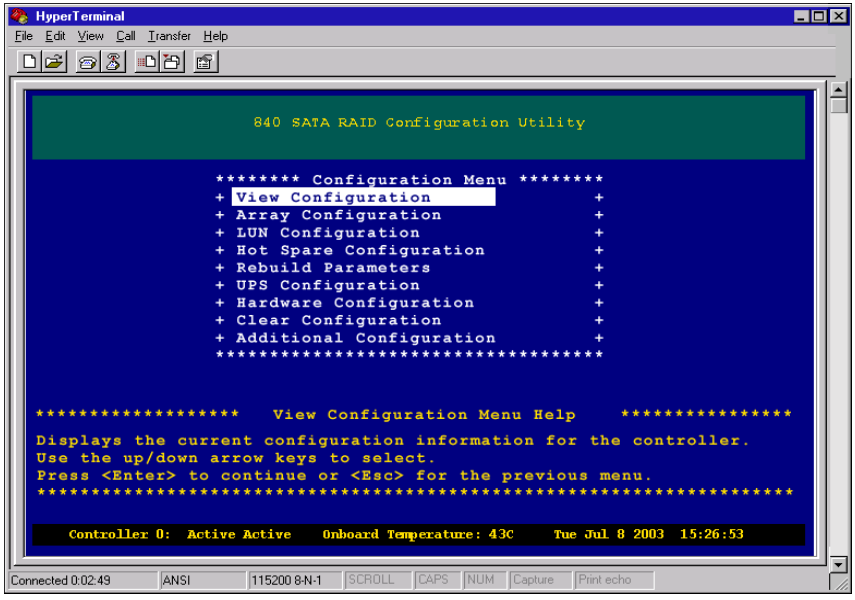

2 From the Configuration Menu, select **Hardware Configuration**, then press ENTER. The *Hardware Configuration Menu* opens.

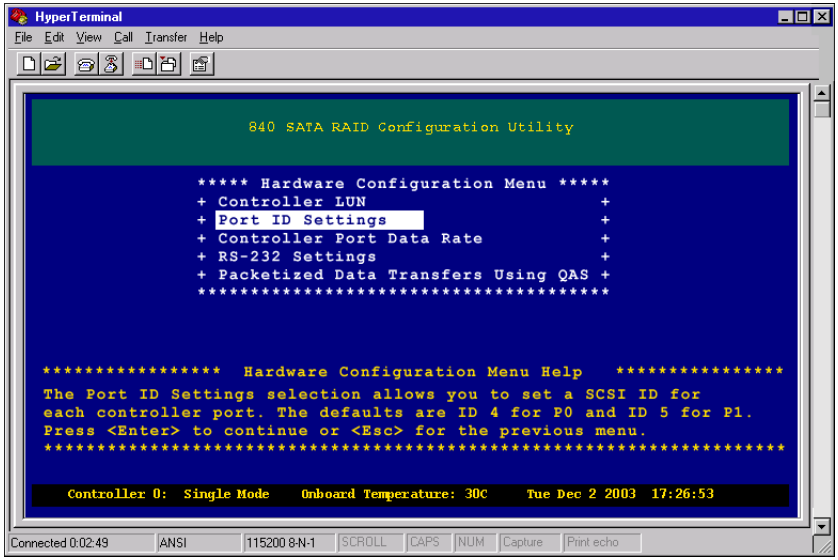

3 From the Hardware Configuration Menu, select **Port ID Settings**, then press ENTER. The *Port ID Menu* opens.

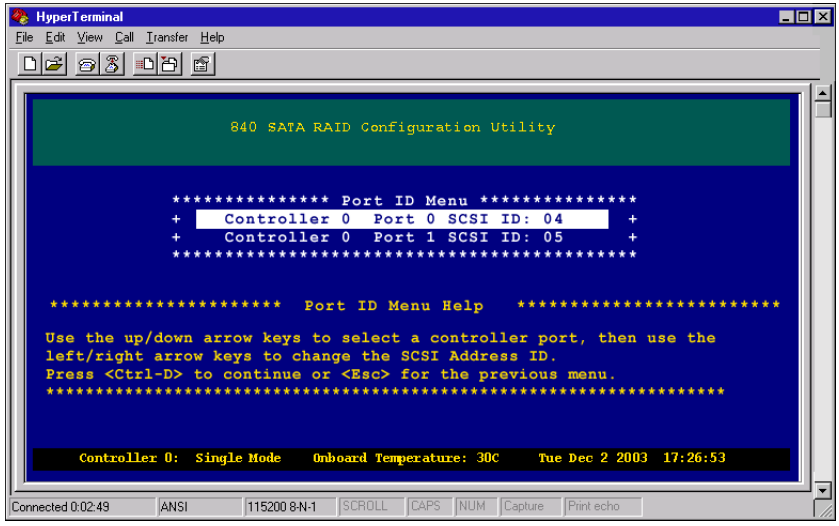

4 Press ENTER to toggle between **Disabled** and **Enabled**.

The Disabled setting allows for soft addressing. The default IDs for the controller ports are Port 0 ID 4 and Port 1 ID 5.

Important

IDs can only be changed when there is a RAID configuration present.

5 Press CTRL + D to save the changes and continue.

6 When you are prompted to confirm the save, choose **Yes**, then press ENTER.

7 Press any key to return to the Main Menu.

Controller port data rate

This option lets you set the data transfer rate of the controller ports.

The settings include 320 MB, 160 MB and 80 MB.

To set the controller port data rate:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

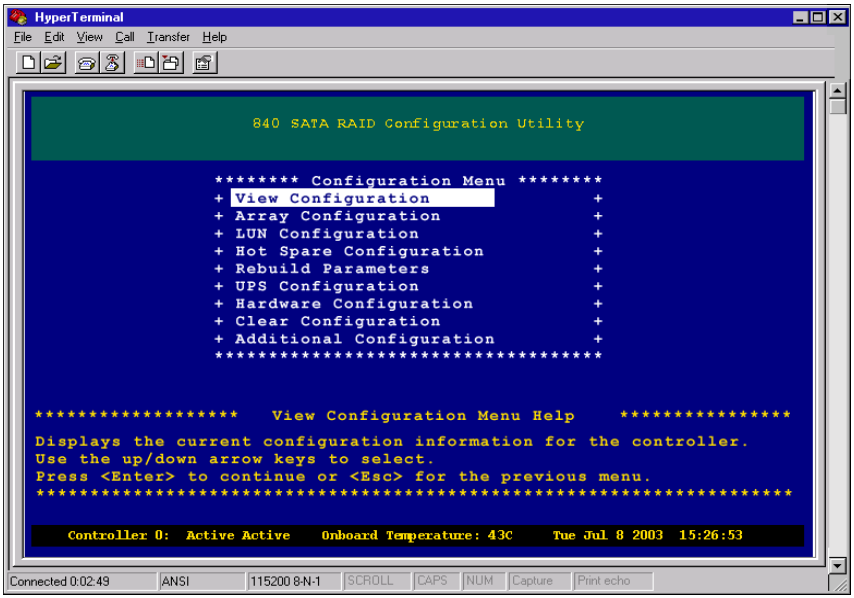

2 From the Configuration Menu, select **Hardware Configuration**, then press ENTER. The *Hardware Configuration Menu* opens.

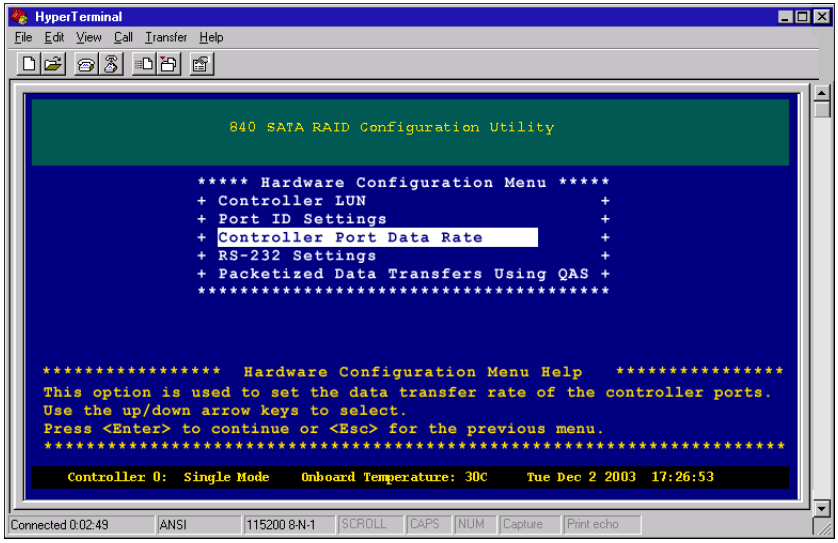

3 From the Hardware Configuration Menu, select **Controller Port Data Rate**, then press ENTER. The *Control Port Data Rate Menu* opens.

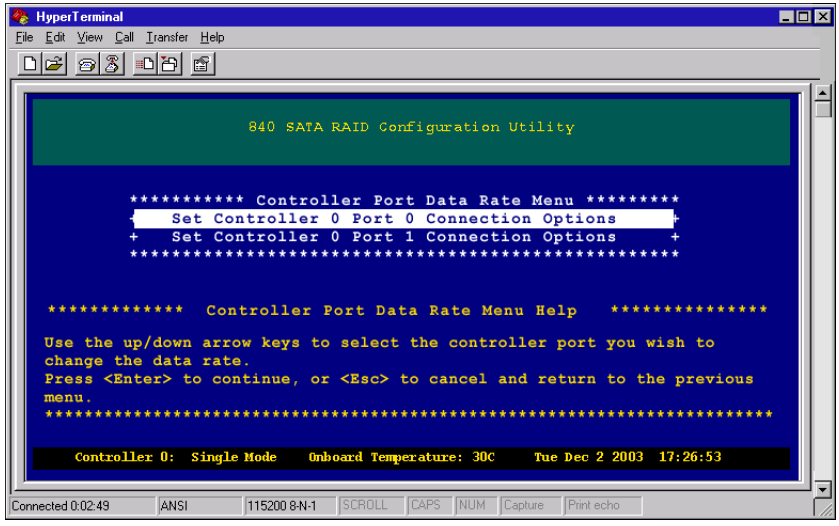

4 Use the up and down arrow keys to choose a controller port, then press ENTER. The *Host Port Data Rate Options screen* opens.

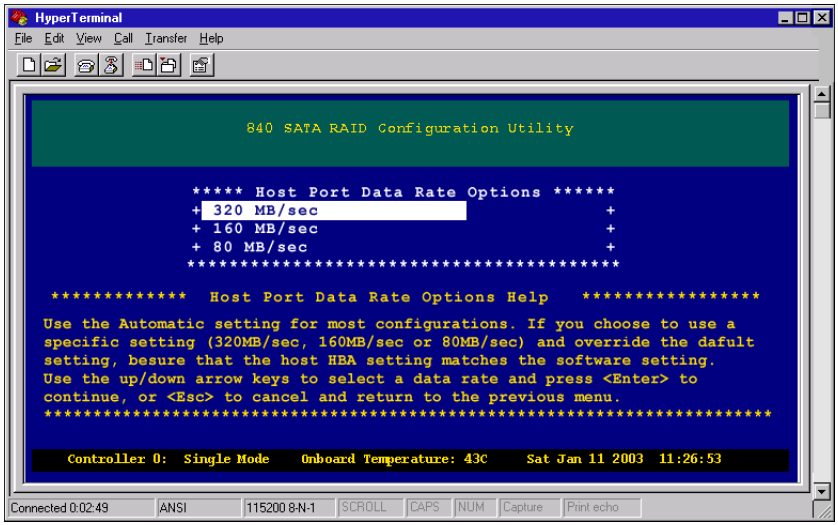

5 Use the up and down arrow keys to select **320MB/sec**, **160MB/sec**, or **80 MB/sec**, then press ENTER.

6 When you are prompted to confirm the change, choose **Yes**, then press ENTER.

7 Press any key to return to the Main Menu.

Setting the RS-232 configuration

The RS-232 Settings option provides you with the option to set the BAUD rate for the controller's RS-232 port.

To change RS-232 settings:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

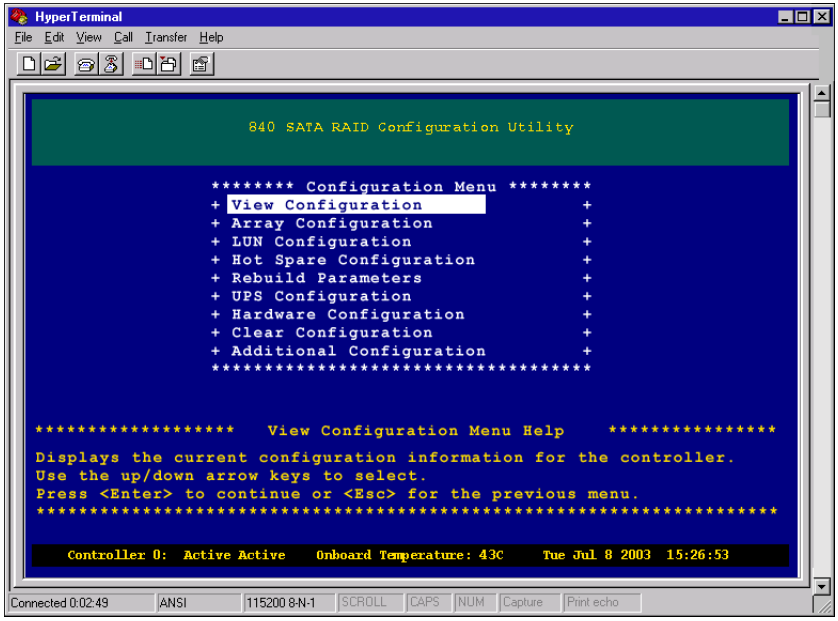

2 From the Configuration Menu, select **Hardware Configuration**, then press ENTER. The *Hardware Configuration Menu* opens.

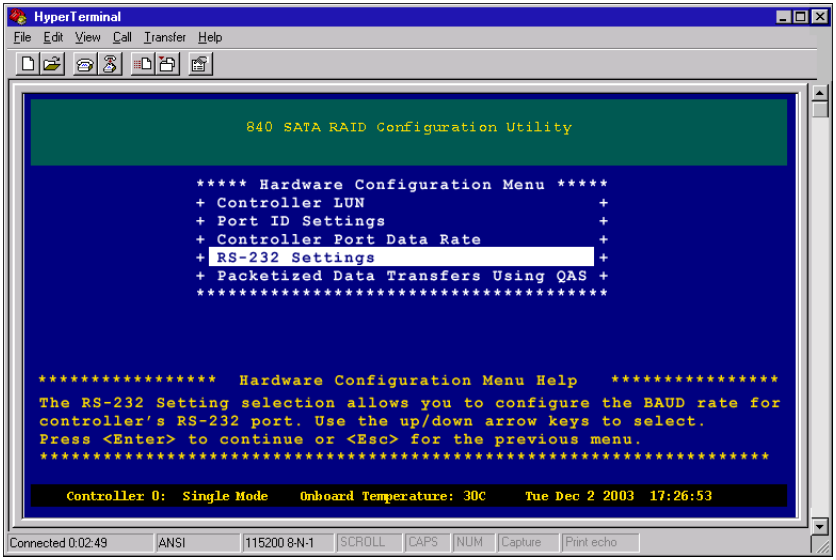

3 From the Hardware Configuration Menu, select **RS-232 Settings**, then press ENTER. The *Baud Rate Selection Menu* opens.

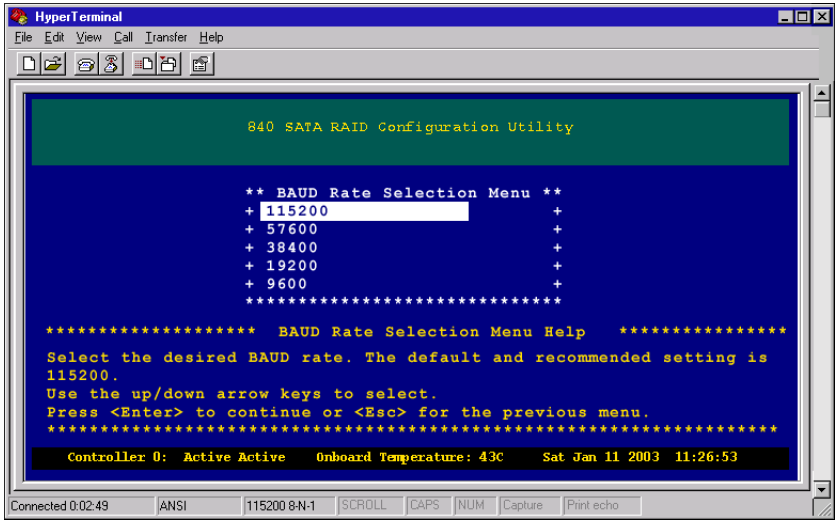

4 Press the up and down arrow keys to select a BAUD rate for the controller RS-232 port (the default and recommended setting is 115,200), then press ENTER.

5 When you are prompted to save the configuration, select **Yes**, then press ENTER.

6 Press any key to return to the Main Menu.

Packetized data transfers/QAS

Packetized protocol is provided for transferring commands, status, and data. The system sends this information in packets that are protected with a CRC. These packets are also sent synchronously to reduce the protocol overhead. This option is set Enabled by default for 320 MB/sec data rates, and is required for 320 MB/sec operations.

Tips & Tricks

The default setting of “Enabled” should be valid for most systems

To set the packetized protocol:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Hardware Configuration Menu* opens.

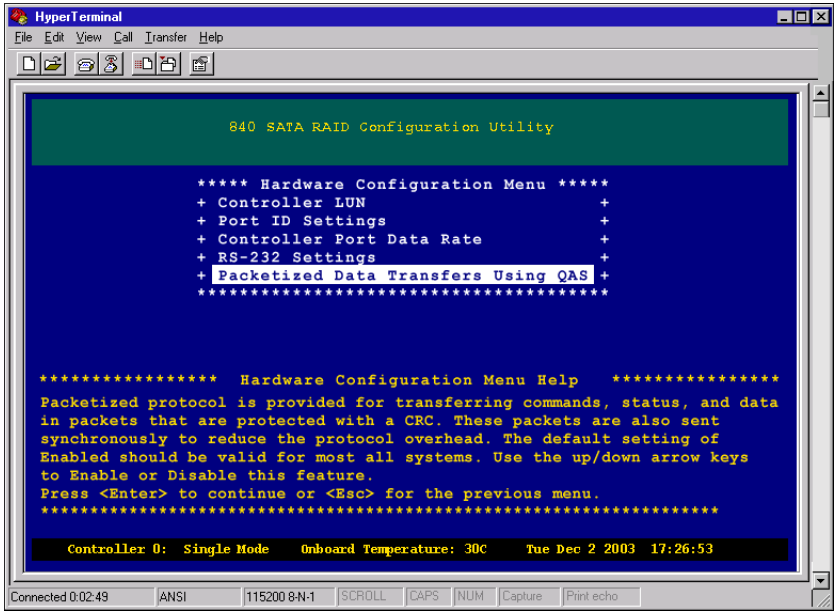

2 Press the up or down arrow keys to select *Packetized Data Transfers Using QAS,* then press **ENTER**. The *Packetized SCSI/QAS* screen opens.

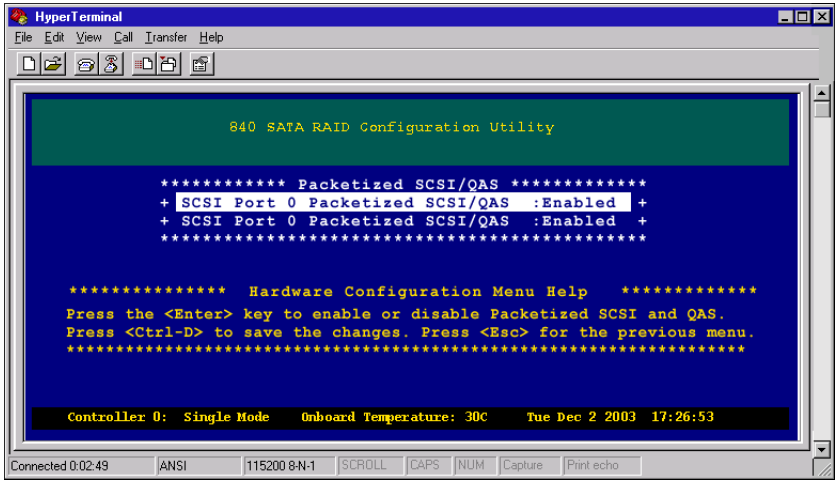

3 Press the up or down arrow keys to select **Enabled** or **Disabled**.

4 Press **CTRL** + **D** to save the changes and continue.

5 When you are prompted to confirm the save, choose **Yes**, then press **ENTER**.

6 Press any key to return to the Main Menu.

Viewing controller information

This option provides the ability to view detailed information on the SATA RAID Controller, such as: Controller WWN, Configuration WWN and Name, Firmware, run-time data and onboard temperature and voltage information.

To view controller information:

1 From the Main Menu, select **Controller Information**, then press ENTER. The *Select Controller Menu* opens.

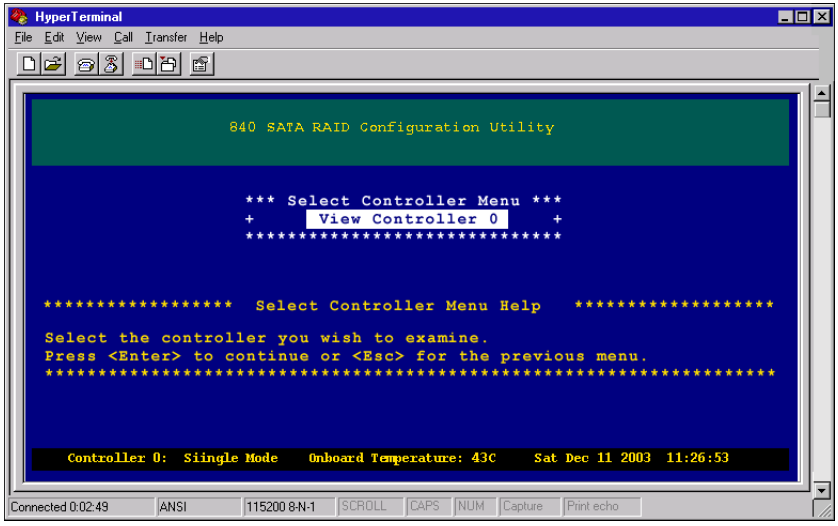

2 From the Controller Information Menu, press the up and down arrow keys to select a Controller, then press **ENTER**. The *Controller Information* screen opens.

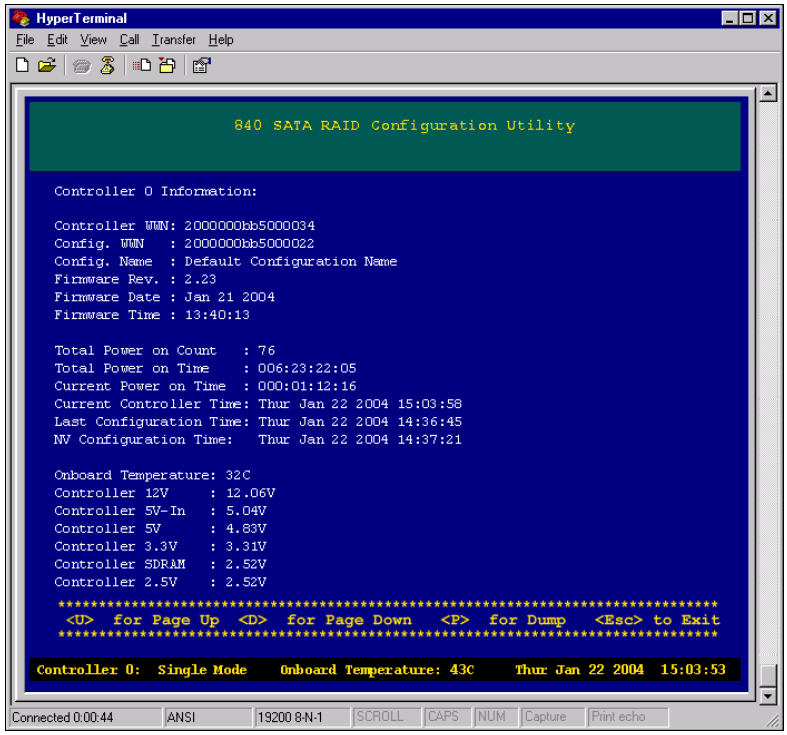

3 Use **U** and **D** to scroll up and down through the controller information displayed.

If you want to capture the information, press **P** to dump the data to a text file.

4 Press **ESC** to return to the previous menu.

Clearing the configuration

To clear the configuration:

1 From the Main Menu, select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

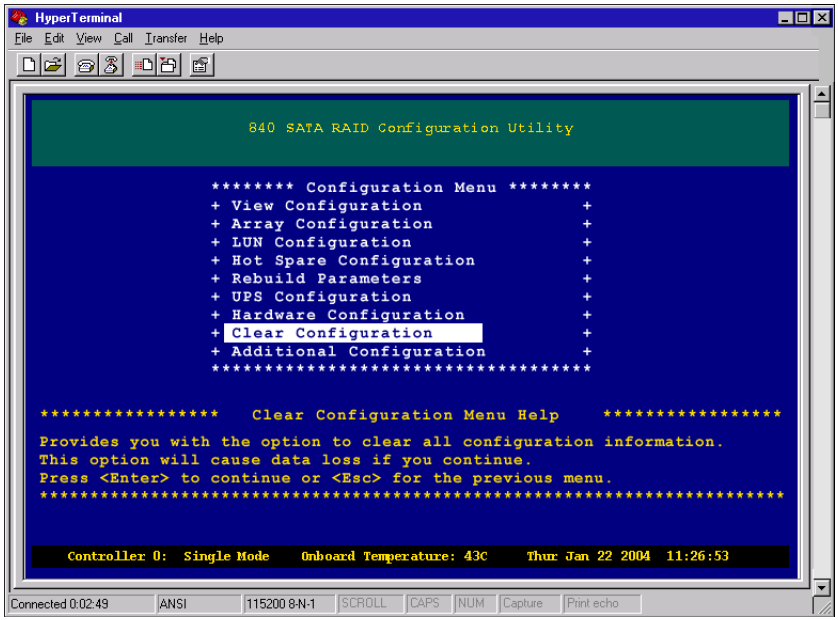

2 From the Configuration Menu, select **Clear Configuration**, then press **ENTER**. The *Clear Configuration Warning* opens.

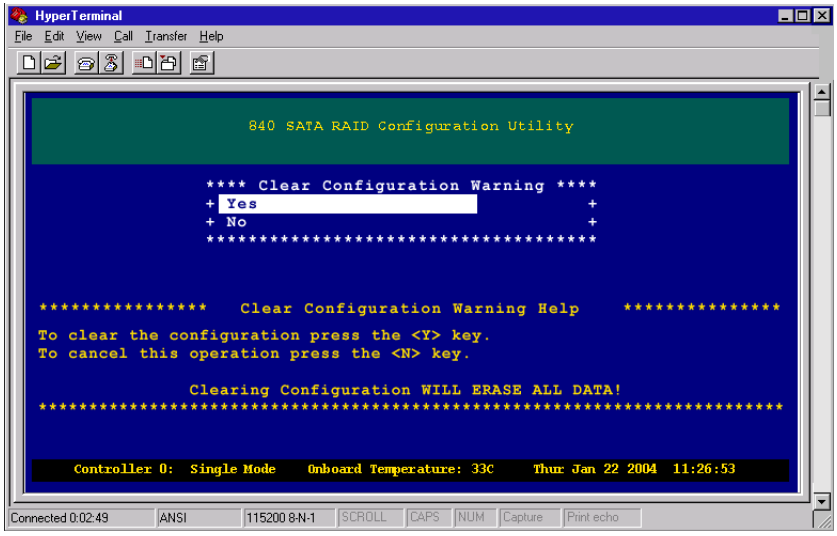

Caution

There is no further warning, clearing the configuration will erase all data.

3 Press **Y** to continue and clear all configuration information or **N** to cancel this operation

4 Press **ESC** to return to the Main Menu.

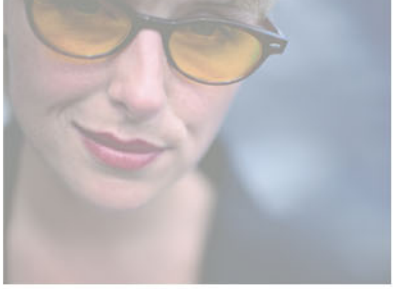

7 Additional Configuration

This chapter provides information on additional configuration. Read this chapter to learn how to:

- Modify configuration names
- Modify configuration WWNs
- Set single controller mode
- Verify background media
- Include enclosure events in the event logs

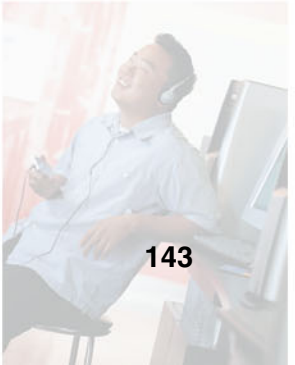

Modifying the configuration name

To modify the configuration name:

1 From the Main Menu select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

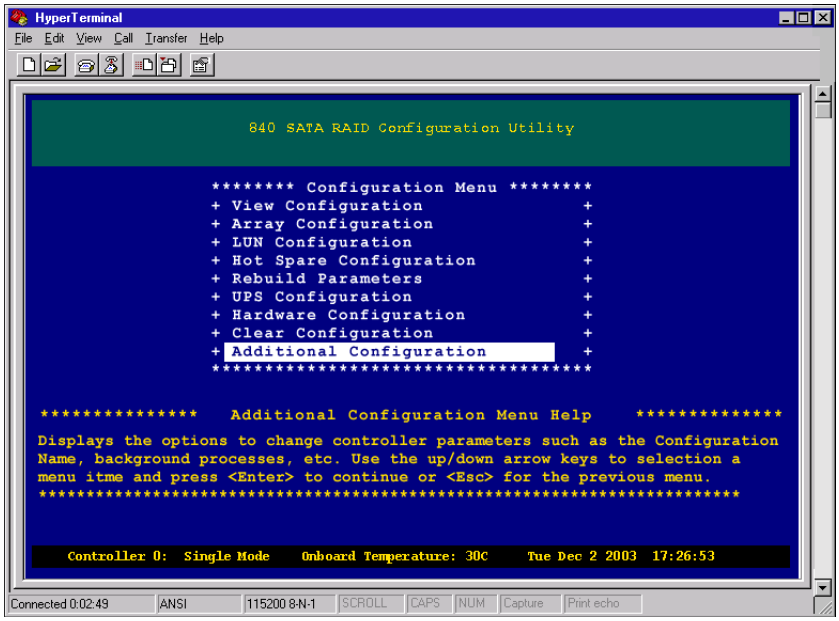

2 From the Configuration menu, select **Additional Configuration**, then press **ENTER**. The *Additional Configuration Menu* opens.

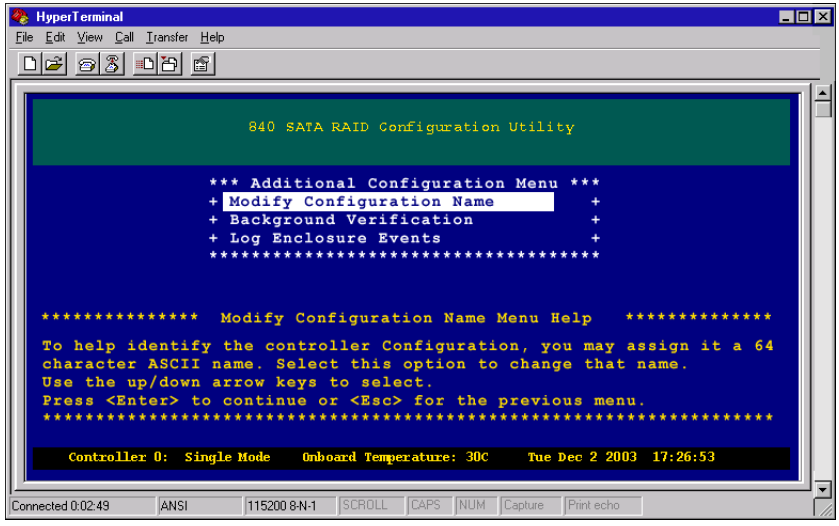

3 Press the up and down arrow keys to select **Modify Configuration Name**, then press **ENTER**. The *Configuration Name* screen opens.

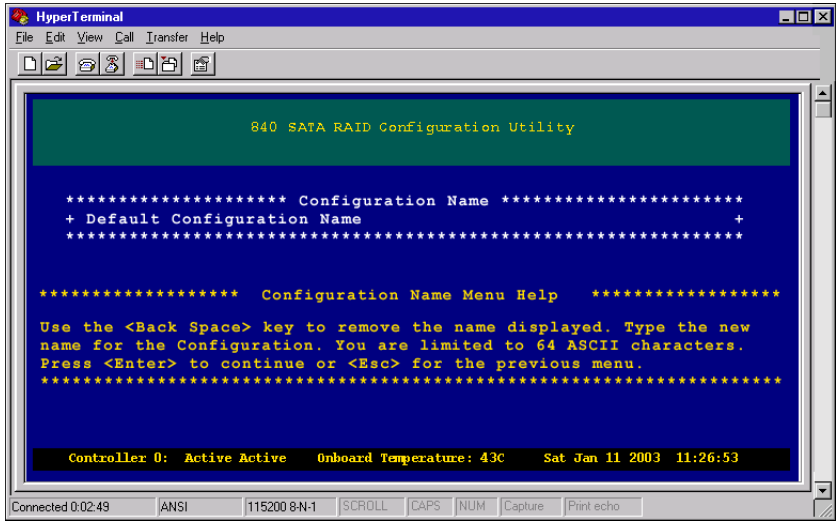

4 Press **BACKSPACE** to delete the existing Configuration Name displayed, type a new name for your controller's Configuration (as many as 64 ASCII characters), then press **ENTER** to continue.

5 When you are prompted to save the configuration, select **Yes**, then press **ENTER**.

6 Press any key to return to the Main Menu.

Verifying background media

This option, when enabled, will automatically verify the media of all drives in the background. If a media error is detected, the controller can automatically re-write the data, providing that the array is in a fault tolerant mode.

To verify background media:

1 From the Main Menu select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

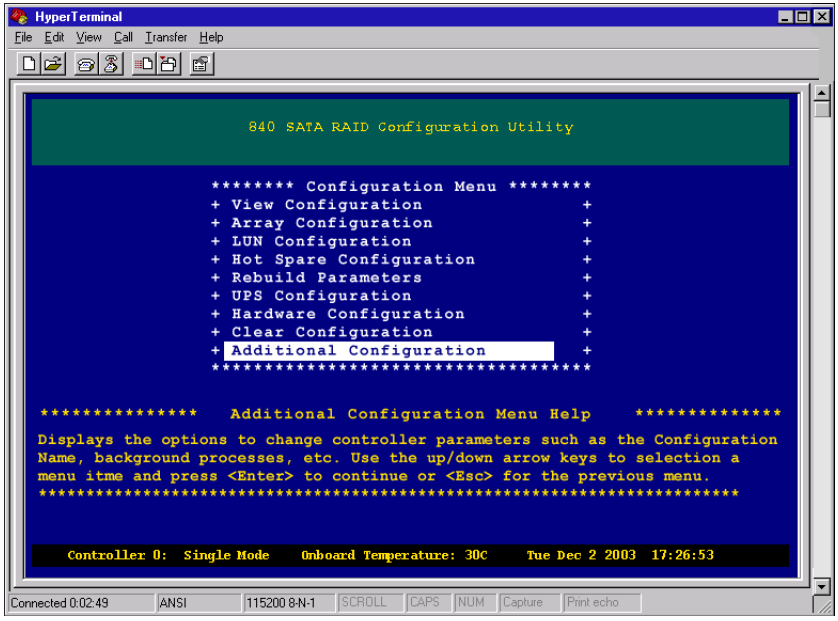

2 From the Configuration menu, select **Additional Configuration**, then press **ENTER**. The *Additional Configuration Menu* opens.

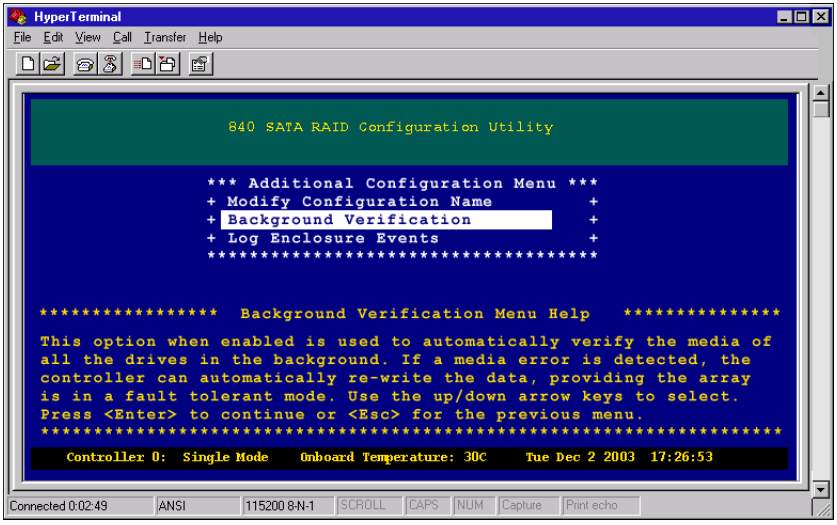

3 From the Additional Configuration Menu, select **Background Verification**, then press **ENTER**. The *Background Verification* screen opens.

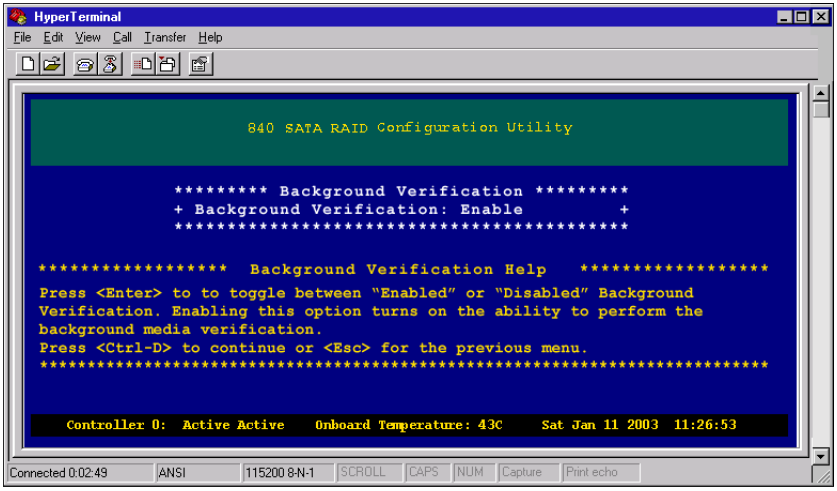

4 Press **ENTER** to toggle between **Enabled** or **Disabled** Background Verification.

5 Press **CTRL** + **D** to continue and save the changes.

6 When you are prompted to save the configuration, select **Yes**, then press ENTER.

7 Press any key to return to the Main Menu.

Including enclosure events in event logs

This option lets you enable or disable whether enclosure events are included in the Event logs. Enclosure events are those which pertain to the enclosure components such as cooling fan status, power supply status, and I/O card status.

To include enclosure events:

1 From the Main Menu select **Configuration Menu**, then press ENTER. The *Configuration Menu* opens.

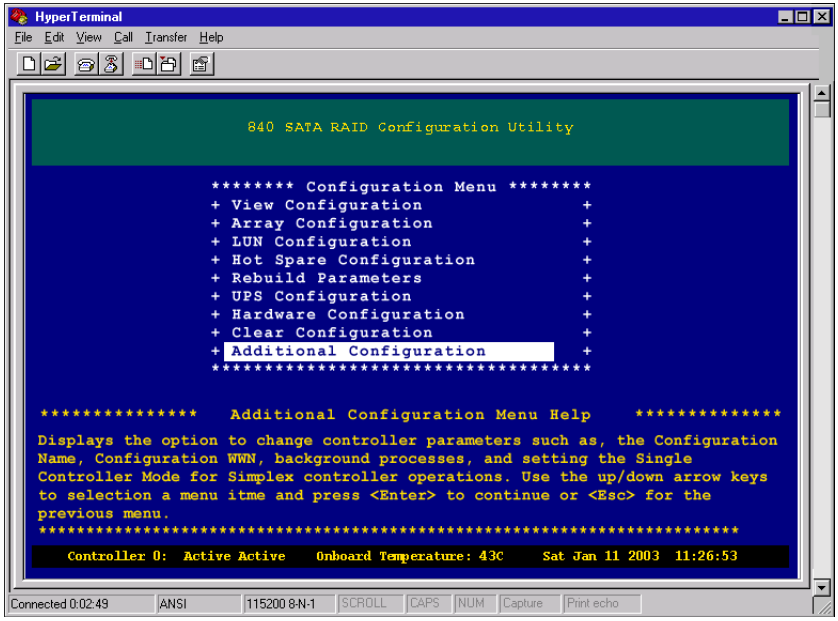

2 From the Configuration menu, select **Additional Configuration**, then press **ENTER**. The *Additional Configuration Menu* opens.

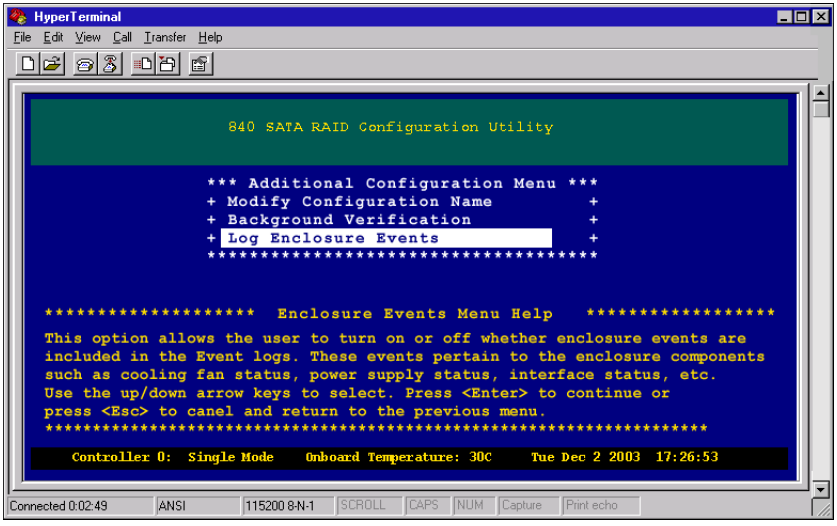

3 From the Additional Configuration Menu, select **Log Enclosure Events**, then press **ENTER**. The *Enclosure Events* screen opens.

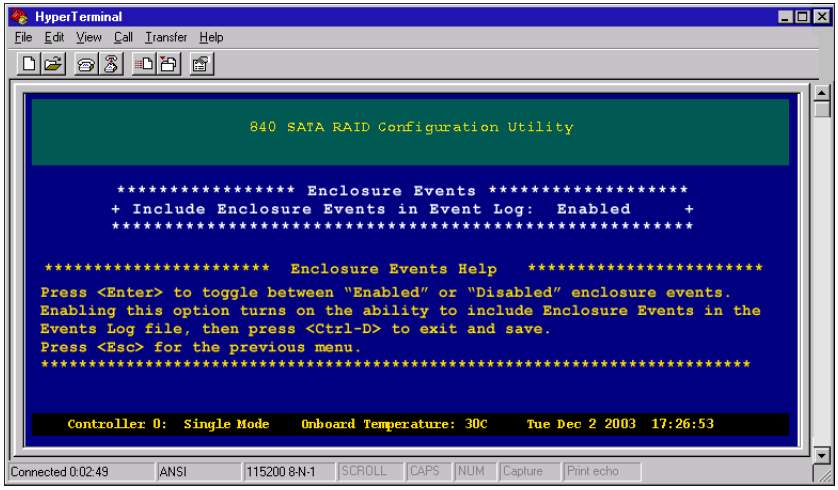

4 Press **ENTER** to toggle between **Enabled** or **Disabled** (SES) enclosure events.

5 Press **CTRL** + **D** to continue and save the changes.

6 When you are prompted to save the configuration, select **Yes**, then press ENTER.

7 Press any key to return to the Main Menu.

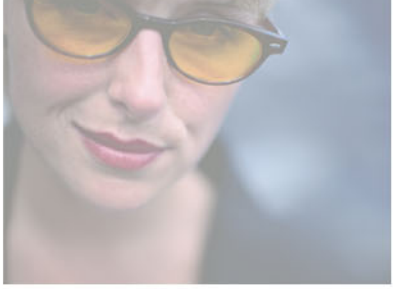

8 Statistics

This chapter provides information on statistics. Read this chapter to learn how to:

- View all statistics
- View access statistics
- View readahead statistics
- View command cluster statistics
- View miscellaneous statistics
- View access statistics on each logical drive
- Clear access statistics on each logical drive
- Clear statistics

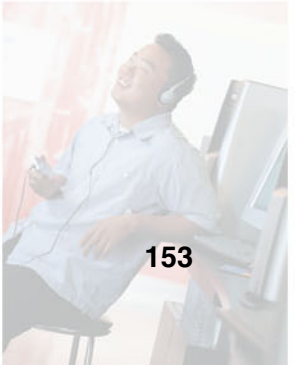

Overview

The RAID controller monitors all incoming commands and calculates various statistics. These statistics are then reported to the controller's serial interface, where you can view this data. The statistics monitored include: Command Count, Command Alignment, Command Size, Readahead Statistics, Write Clustering Statistics, and RAID 5/50 Write Statistics.

The controller maintains individual access statistics for each LUN and controller port. These can be useful to help balance the load from the host. Identical statistics are maintained for both reads and writes.

Viewing all statistics

To view all statistics:

1 From the Main Menu, select **Statistics Menu**, then press ENTER. The *Select Controller Statistics Menu* opens.

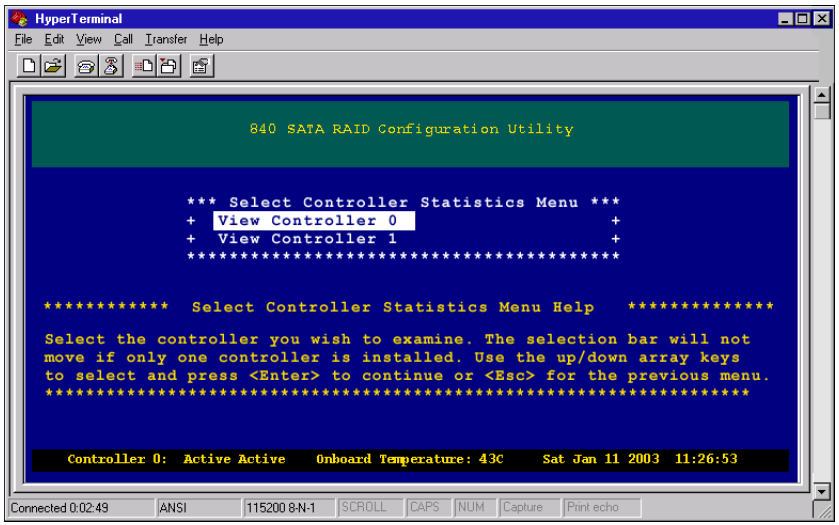

2 From the Select Controller Statistics Menu, select the controller on which you want to view the statistics, then press ENTER. The *Statistics Types* screen opens.

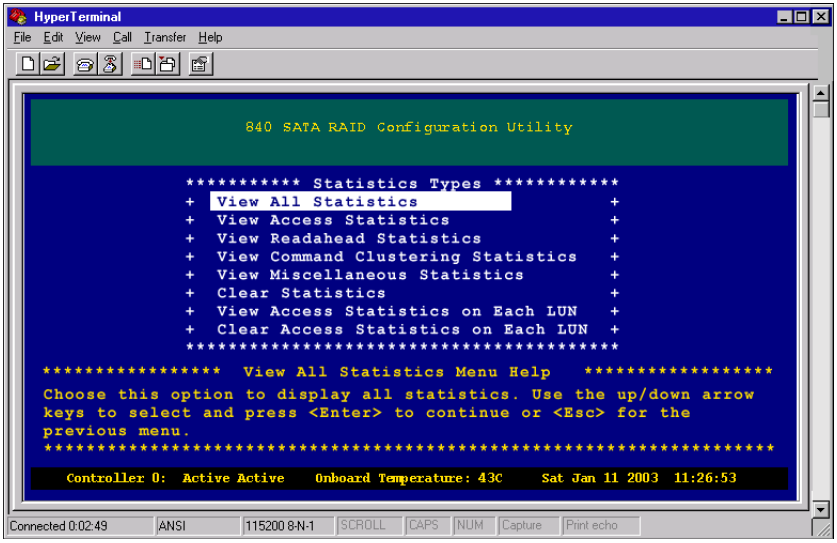

3 From the Statistics Types screen, select **View All Statistics**, then press ENTER. The *All Statistics* screen opens.

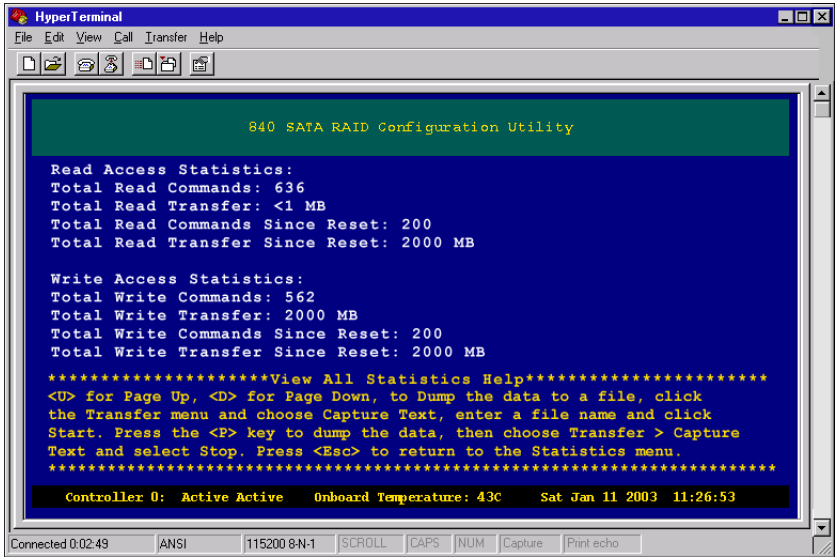

4 Press **U** to page up, **D** to page down, and **P** to dump the information to a file.

The balance of this chapter covers each statistics group selected, with an explanation of each item.

Dumping data to a text file

The dump function enables you to write the data to a text file.

To dump data to a text file:

1 Click the **Transfer** list and select **Capture Text**.

2 When a message prompts you to name the file, type a file name and destination or choose the default name, then click **Start**.

3 Press the **P** to dump the data.

Tips & Tricks

You will notice the cursor jumps to the lower-right-corner of the screen and after completion will return to the upper-left-corner of the screen.

4 Click the **Transfer** list, then select **Capture Text** > **Stop**.

5 Access the text file using a text editor.

Viewing access statistics

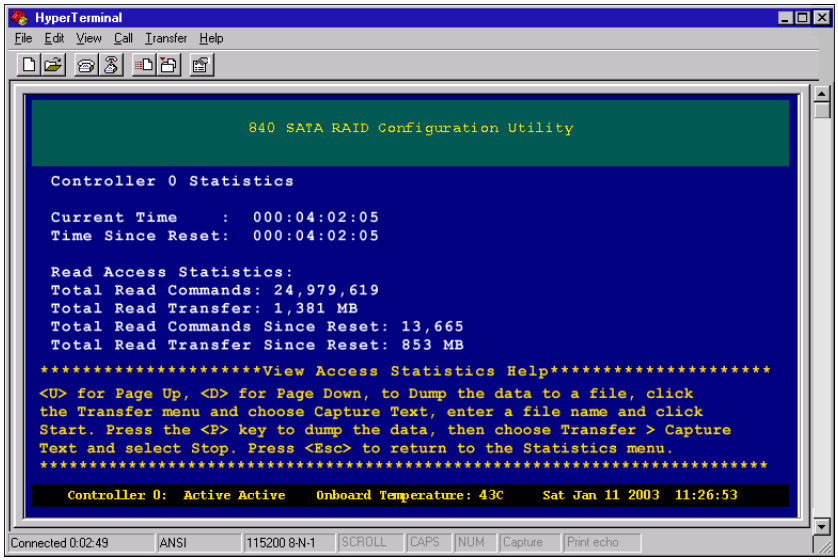

The statistics groups can be viewed all together or individually, as well as on a LUN (logical drive) or Port basis.

The *View Access Statistics* screen is accessed by going to the *Statistics Types* menu (see "Viewing all statistics" on page 155), then selecting **View Access Statistics**.

Each time statistics are viewed, the controller first outputs the current time and the time since it was last reset. The statistics can be reset at any time. This is useful in determining the access pattern for a particular test or period of time.

Statistic	Description
Total Read Commands	This is a count of the total number of read accesses (for all ports and logical drives) that have occurred since the controller was first powered on. This value is never reset and indicates the number of commands that were sent to the controller from the time it was manufactured.
Total Read Transfer	This is a count of the total number of 512-byte blocks (for all ports and logical drives) that have been read since the controller was first powered on. This value is never reset and indicates the amount of data transferred by the controller from the time it was manufactured.

Statistic	Description
Total Read Commands Since Reset	This is a count of the total number of read accesses that have occurred since the statistics were reset or the controller was last powered on.
Total Read Transfer Since Reset	This is a count of the total number of 512-byte blocks that have been read since the statistics were reset or the controller was last powered on.
Alignment	This is the percentage of commands whose address is aligned on the specified address boundary. The alignment of a command from a host system is determined by the command's address. In an optimal system, a write of one chunk of data would reside exactly within a chunk on one disk. However, if this is not the case, this write will be split up into two separate writes to two different data drives. This of course will have a negative effect on performance. To overcome these problems, you can, with more sophisticated operating systems, set the access size and alignment to an optimal value. These statistics can help you to tune the operating system.

Statistic	Description
Access Size	This is the percentage of commands of the size specified. On occasions, an application may specify a particular access size, but the operating system can reduce this. This can lead to performance problems. For example, on Windows NT, Windows 2000, and Windows 2003 the maximum command size that the operating system will allow is 1MB. Additionally, some HBAs further reduce this value by default. However, some benchmark programs allow a larger access size to be specified, resulting in misleading results. By viewing the statistics on the controller, it is possible to determine the exact size of commands sent by the host, and so determine whether the results are appropriate. The access size, in conjunction with the alignment, gives an indication of how many drives are involved in an access. For example, consider a RAID 5/50 array with a chunk size of 64 K. In this case, a 64 K access with an alignment of 8 K will actually involve 2 data drives, since it needs to access some data in the first drive, and the remaining data in the next drive. This is clearly inefficient, and could be improved by setting the alignment to 64 K on the operating system. If that is not possible, using a larger chunk size can help, as this reduces the number of accesses that span chunks. Aligning an access on the same value as the access size will improve performance, as it will ensure that there are not multi-chunk accesses for commands that are smaller than a chunk size.

Viewing readahead statistics

If sequential read commands are sent to the controller, it assumes that the commands that follow can also be sequential. It can then go and perform a read of the data before the host requests it. This improves performance, particularly for smaller reads. The size of the readahead is calculated based on the original command size, so the controller does not read too much data. The controller maintains statistics for all readahead commands performed.

The *View Readahead Statistics* screen is accessed by going to the *Statistics Types* menu (see "Viewing all statistics" on page 155), then selecting **View Readahead Statistics**.

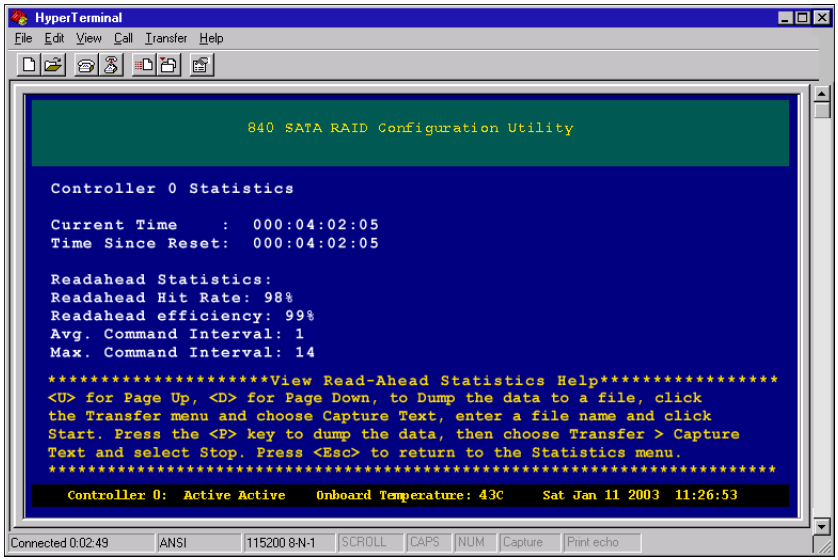

Statistic	Description
Readahead Command Hit Rate	This is the percentage of read command hits versus the total number of read commands that have been issued. This gives an indication of the sequential nature of the data access pattern from the host.
Readahead Command Efficiency	This is the percentage of the number of read command hits versus the projected number of readahead command hits. This is a measure of the efficiency of the readahead algorithm. A low value means that much of the data that the controller reads in the readahead command is not actually requested by the host.

Statistic	Description
Sequential Command Interval	In determining whether to perform a readahead or not, the controller will search back in the command queue whenever it receives a new read command that is not satisfied by an existing readahead cache buffer. In a multi threaded operating system, commands from one thread can be interspersed with commands from another thread. This requires that the controller not just check the immediately previous command. The controller will search back for a number of commands, to see if the new command is exactly sequential to any one of these previous commands. If it is, then the controller determines that the data access pattern is sequential, and so performs a readahead. These statistics record the average number of commands the controller must search back for when it finds a sequential command match, the maximum number, and also the percentage for each one of these values. These give an indication of the multi threaded nature of the host.

Viewing command cluster statistics

To increase performance, the controller can cluster sequential write commands together to create a larger write command. This results in less commands being sent to the hard drives. Additionally, if sufficient data is clustered by the controller, it can perform a a full stripe write for RAID5/50 arrays. This significantly improves performance. In cases where the host does not send a sufficient number of outstanding writes, writeback cache can be used to delay the write to disk, increasing the likelihood of clustering more data.

The *View Command Clustering Statistics* screen is accessed by going to the *Statistics Types* menu (see "Viewing all statistics" on page 155), then selecting **View Command Clustering Statistics**.

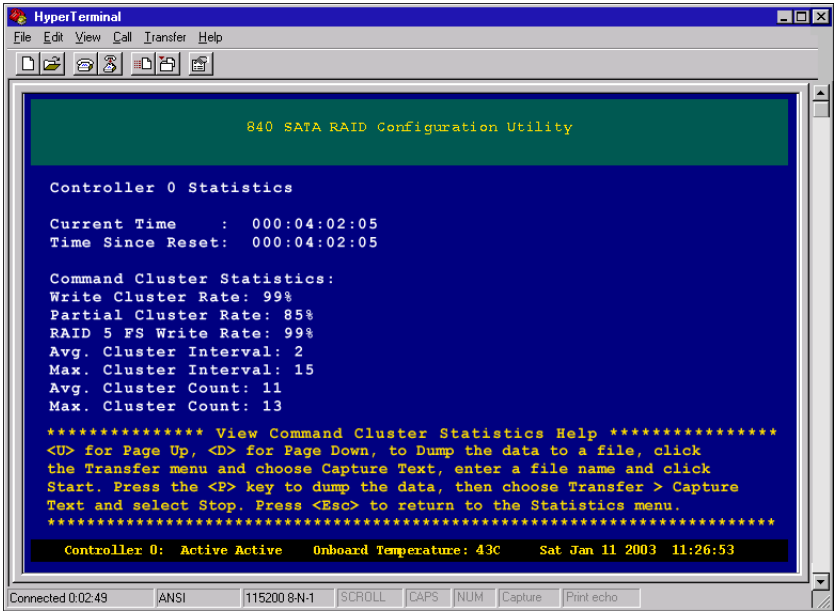

Statistic	Description
Write Cluster Rate	This is the percentage of the number of write commands that are part of a cluster versus the total number of write commands that have been issued. This gives an indication of the sequential nature of the data access pattern from the host and of the performance of the writeback cache.
RAID 5/50 Partial Command Cluster Rate	This is the percentage of the number of clustered commands that wrap around the end of a stripe versus the total number of clustered commands. This gives an indication of the alignment of the operating system.
RAID 5/50 Full Stripe Write Rate	This is the percentage of the amount of data that is written as a full stripe write versus the total amount of data written. This gives an indication of the sequential nature of the data access pattern from the host and of the performance of the writeback cache, for RAID 5/50 drive ranks.
Command Cluster Interval	In determining whether to cluster write commands, the controller will search back in the command queue whenever it receives a new write command. In a multi threaded operating system, commands from each thread can be interspersed with commands from another thread. This requires that the controller not just check the immediately previous command. The controller will search back for a number of commands, to try to determine if the new command is exactly sequential to any one of these previous commands. If it is, then the controller determines that it can cluster these commands. These statistics record the average and maximum number of commands the controller must search back for when it finds a sequential command match and the percentage for each one of these values.
Command Cluster Count	When the controller clusters a write command, it can cluster a large number of them together. These statistics record the average and maximum number of commands the controller clusters and the percentage for each one of these values.

Viewing miscellaneous statistics

To aid in tuning performance, some additional statistics have been included for development use under the **View Miscellaneous Statistics** selection.

The *View Miscellaneous Statistics* screen is accessed by going to the *Statistics Types* menu (see "Viewing all statistics" on page 155), then selecting **View Miscellaneous Statistics**.

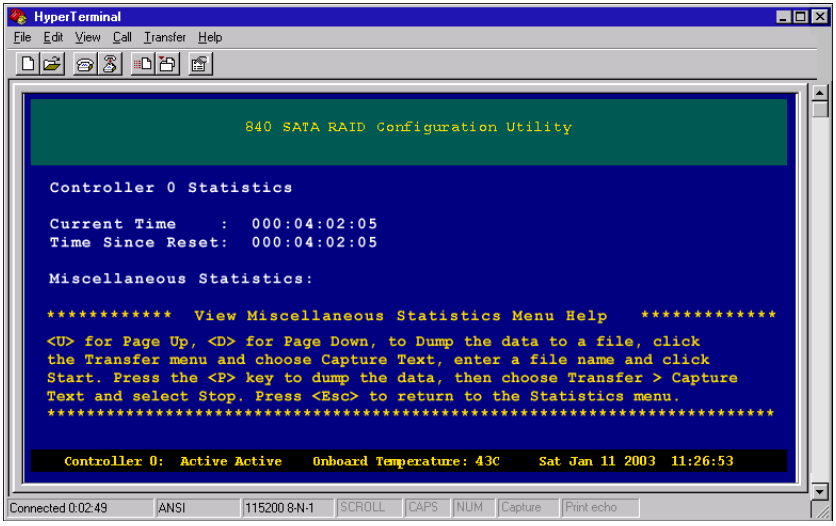

Viewing access statistics on each logical drive

To view access statistics for each logical drive.

1 From the Main Menu, select **Statistics Menu**, then press ENTER. The *Select Controller Statistics Menu* opens.

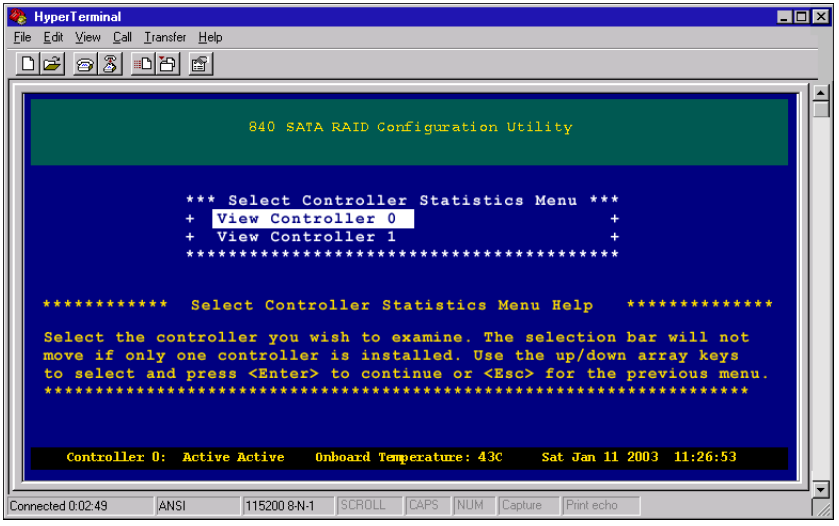

2 From the Select Controller Statistics Menu, select the controller on which you want to view the statistics, then press **ENTER**. The *Statistics Types* screen opens.

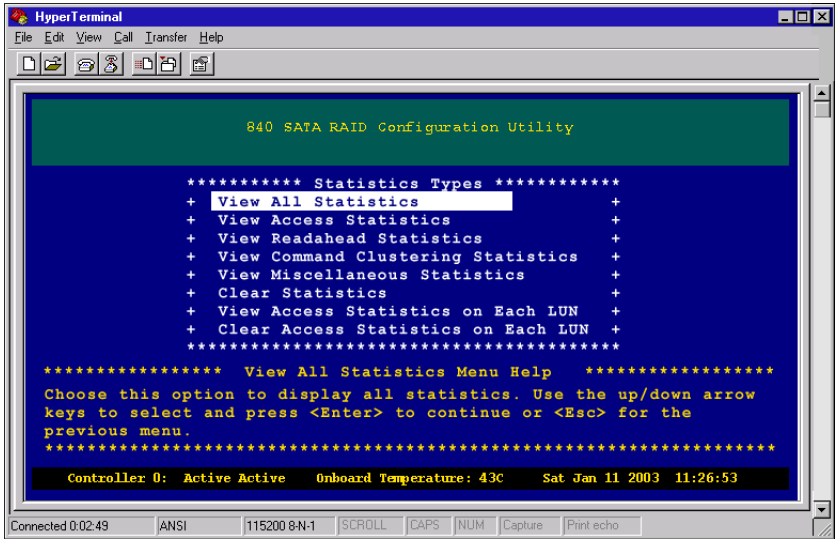

3 From the Statistics Types screen, select **View Access Statistics on Each LUN**, then press **ENTER**. The *Select LUN for Host LUN Statistics* screen opens.

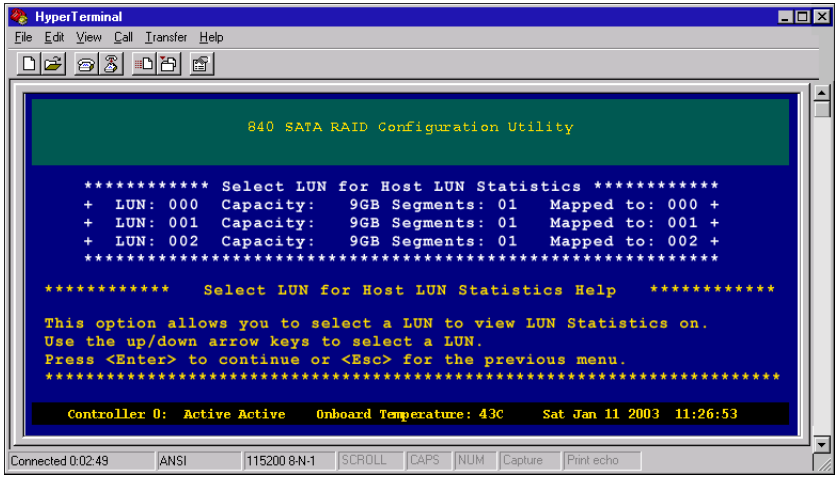

4 From the Select LUN for Host LUN Statistics screen, select a LUN on which to view statistics, then press **ENTER**. The *Select a Port Access to the Host LUN* screen opens.

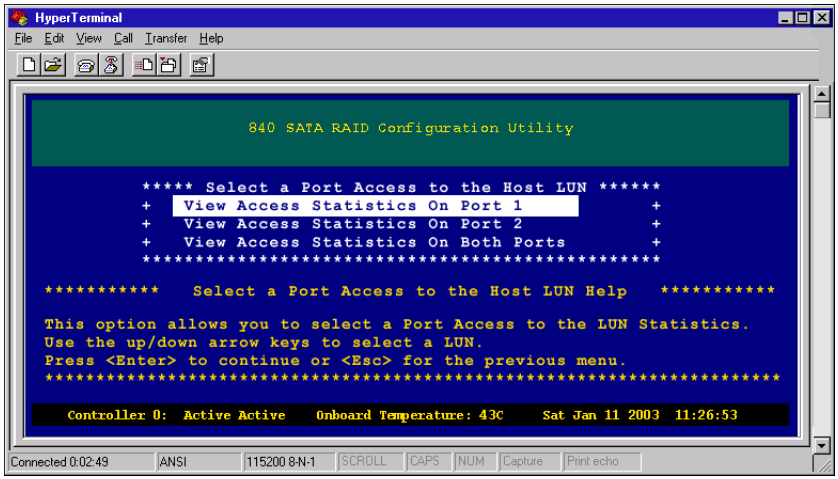

5 Select a Controller Port on which to view Statistics (either Port 1, Port 2 or Both Ports), then press **ENTER**. The *View Access Statistics on Each LUN* screen opens.

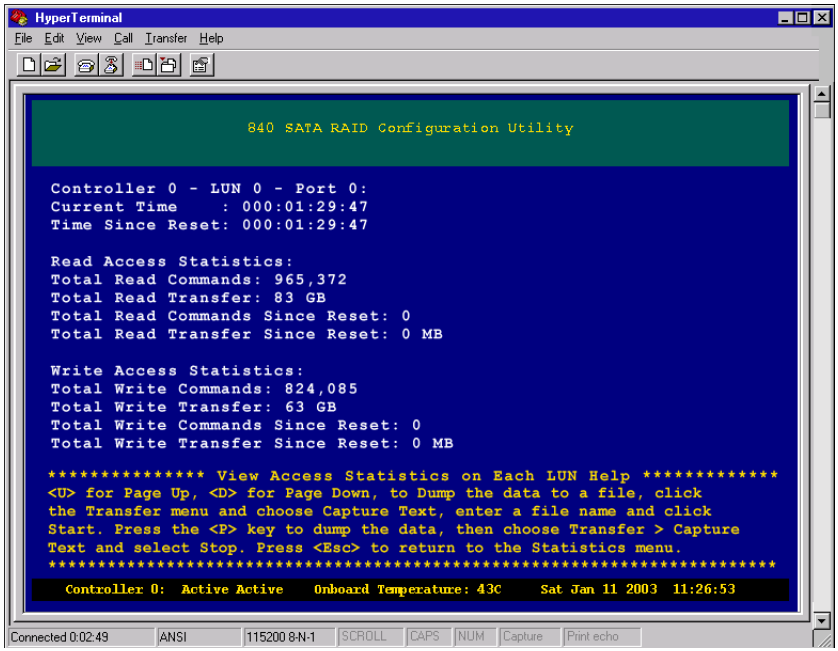

Clearing access statistics on each logical drive (LUN)

To clear the access statistics for each logical drive:

1. Access the *View Access Statistics on Each LUN* screen (see "Viewing access statistics on each logical drive" on page 166).
2. Select a logical drive on which to clear statistics, then press **ENTER**.
3. Select a Controller Port on which to clear Statistics (either Port 1, Port 2 or Both Ports), then press **ENTER**.
4. Select **Yes** to clear the statistics.

Clearing all statistics

To clear statistics:

1. From the Main Menu, select the **Statistics Menu**, then press **ENTER**. The *Select Controller Statistics Menu* opens.

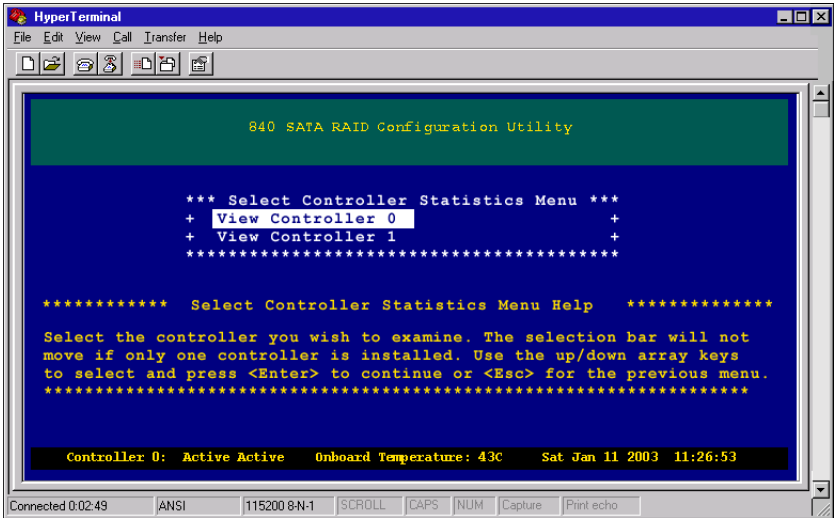

2 From the Select Controller Statistics Menu, select the controller on which you want to clear the statistics, then press **ENTER**. The *Statistics Types* screen opens.

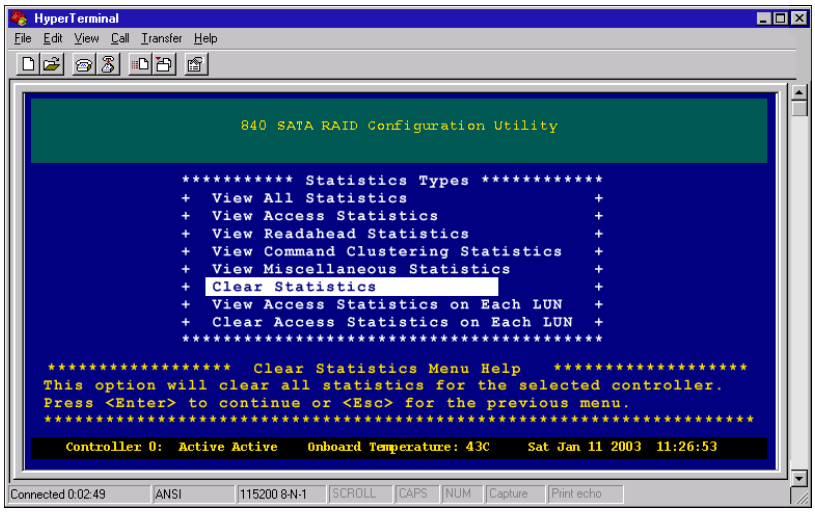

3 From the Statistics Types screen, select **Clear Statistics**, then press **ENTER**. The *Clear Statistics* screen opens.

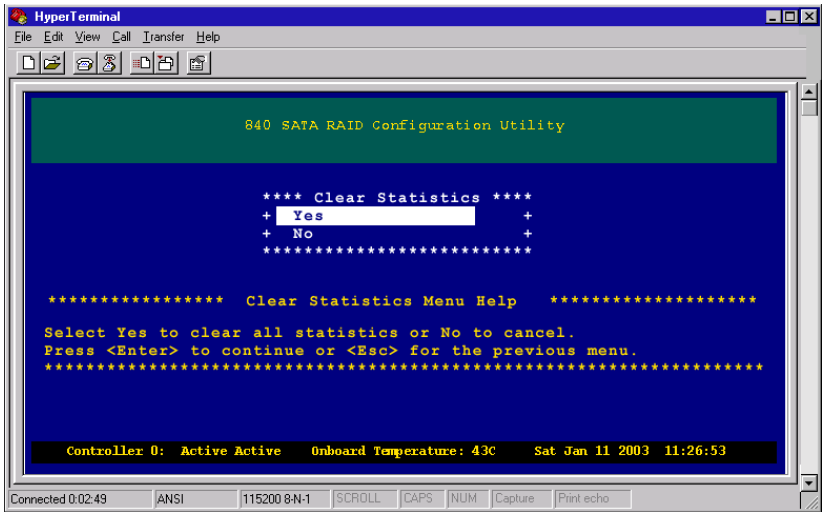

4 Select **Yes** to clear all statistics, then press **ENTER** to continue.

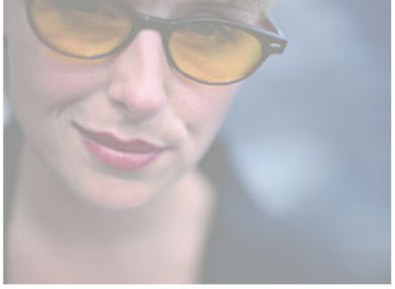

9 Event Logs

This chapter provides information on event logs. Read this chapter to learn how to:

- Access event logs
- Understand controller events
- Understand drive events
- Understand controller drive port (host) events
- Understand enclosure events
- Understand failed drives
- Clear event logs

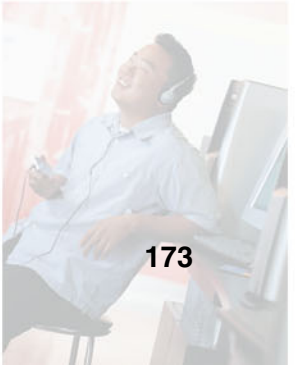

Accessing event logs

The RAID controller has a comprehensive, non-volatile event log that can be used for monitoring controller events and error conditions. The following event types are logged: Drive Failures and Errors, Controller Failures and Errors, (SES) Enclosure, Temperature and Voltage Errors.

The maximum event log size is 4096 entries, and older events are overwritten as necessary. Some repetitive events are appended to previous events, so entries are not used up unnecessarily. A time stamp accompanies each event, so you can easily determine when the event happened. The time stamp of both controllers is synchronized to make the troubleshooting tasks easier.

To access event logs:

1 From the Main Menu, select **Event Logs Menu**, then press ENTER. The *Select Controller Event Log Menu* opens.

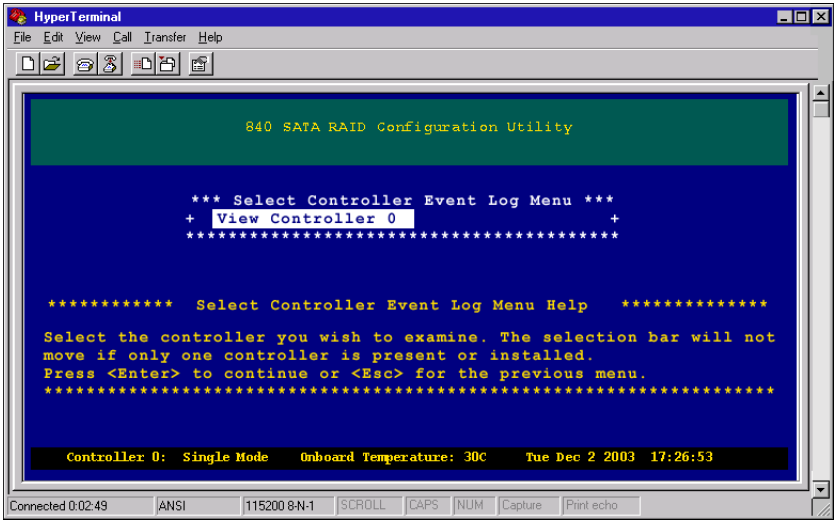

2 From the Select Controller Event Log Menu, select the controller for which you want to view the events, then press **ENTER**. The *Event Log Types* menu opens.

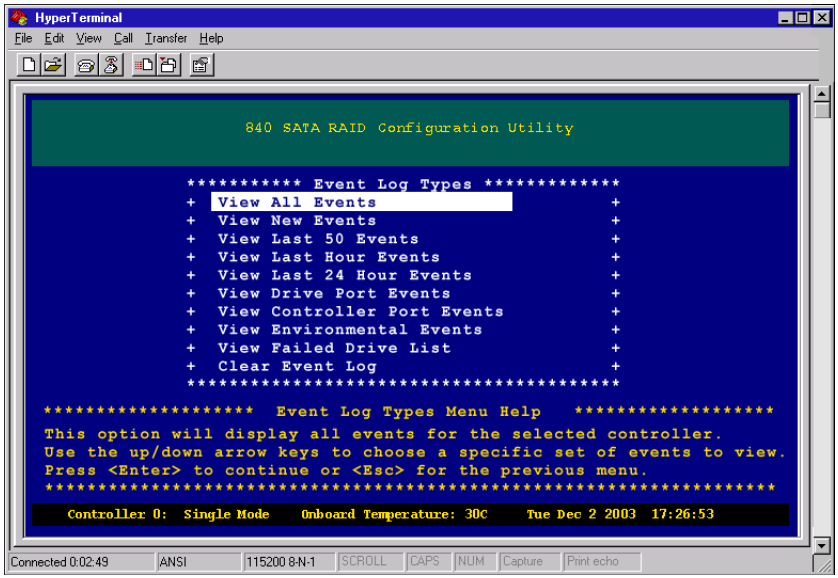

3 From the Event Log Types menu, select a specific set of events you want to view, then press ENTER. For example, if you select **View All Events,** the *View All Events* screen opens.

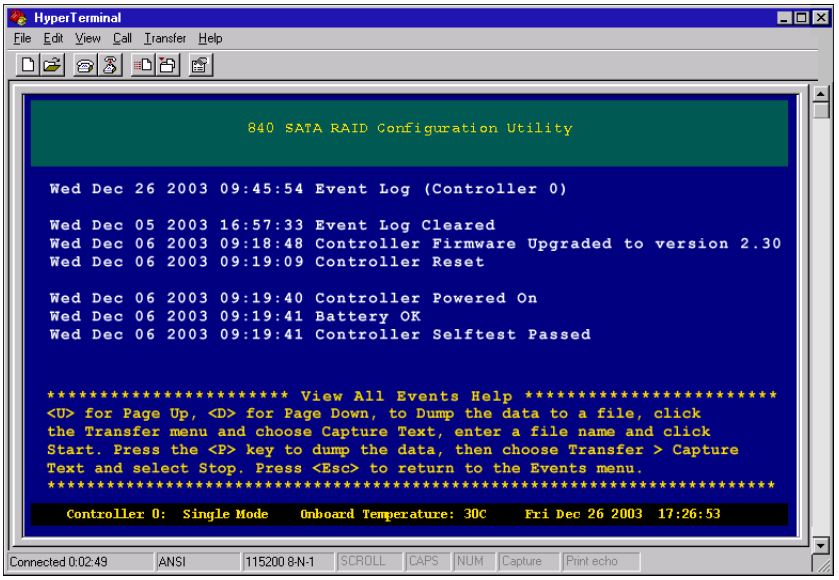

The balance of this chapter covers each Event Log group available, with an explanation of each item.

The events can be viewed all together or for a specific time period or device.

First, the current time and controller number are displayed, then the events are listed. You can view a certain type of event, all events, or all events in a certain time frame. We recommend that you capture the output to a log file, since the number of events can exceed the terminal program's buffer.

These options are as follows:

Event Log	Description
View All Events	This option lists all events. The number of events is indicated.
View New Events	This lists all events that have occurred since the event log has last been viewed.

Event Log	Description
View Last 50 Events	This option lists the last 50 events.
View Last Hour/24 Hour Events	This lists the events that occurred in the specified time.
View Environmental Events	This lists events such as temperature warnings, voltage level errors, battery errors, etc.
View Failed Drive List	This is a list of drives that the controller has marked as having failed. These drives have either failed in operation, or been removed.

Controller events

Following is a brief description of events that relate to the controller and configuration.

Controller Event Messages	Description	Cause	Action
Array *xx* Cache Disabled [Reasons]	The controller has disabled the writeback cache for the indicated reasons, as follows: In Preparation For Shutdown: The controller has received a command to disable the writeback cache in preparation for shutdown. Battery - The controller has disabled the cache because the battery is either not present or fully charged. Critical - The controller has disabled the cache because the array has become critical. UPS - The controller has disabled the cache because a configured UPS has indicated power has been removed or communications has been lost with the UPS.	Disabling of writeback cache for the indicated reasons: ▪ A “prepare for shutdown” command has been received by the controller. ▪ The battery is not charged or present. ▪ The array has become critical. ▪ The UPS has failed.	▪ No action necessary. ▪ Charge the backup battery or re-install the battery. ▪ Resolve the array issue and rebuild the array. ▪ Replace the failed UPS.
Array *xx* Cache Re-Enabled	The controller has re-enabled the writeback cache.	Re-enabling of writeback cache.	No action necessary.
Battery Failure	This indicates that the onboard cache backup battery (if present) has failed.	Battery failure.	Replace battery in the controller.
Battery OK	This indicates that the onboard cache backup battery (if present) is fully charged.	Battery charging complete.	No action necessary.

Controller Event Messages	Description	Cause	Action
Battery Charging	This indicates that the onboard cache backup battery is charging.	Battery charging started.	No action necessary.
Cached Data Lost *xxx* Entries, *yyy* MB	If the controller is powered off with writeback cache still present, it is necessary to flush this data when power is restored. If power is removed for an extended time, the onboard battery will deplete, and the cached data will be lost.	Failure of power for an extended time with writeback cache present.	Check the file system.
Cntrl *x* Temp Exceeded: *yy*	The current internal controller temperature is indicated in degrees Celsius. The controller's internal temperature has exceeded the allowable limit, and the controller will shut down to prevent damage to internal components.	Blocked fan. Failing fan. Elevated ambient temperature.	Check enclosure for sufficient air flow. Check for a failed fan, if found replace cooling fan module. Check the ambient temperature of the environment, decrease the local ambient temperature.
Cntrl *x* Temp Warning: *yy* Shutdown will occur at zzC	The current internal controller temperature is indicated in degrees Celsius. The controller's internal temperature is approaching the maximum allowable limit.	Blocked fan. Failing fan. Elevated ambient temperature.	Check enclosure for sufficient air flow. Check for a failed fan, if found replace cooling fan module. Check the ambient temperature of the environment, decrease the local ambient temperature.

Controller Event Messages	Description	Cause	Action
Controller *x* Failed/Removed	The specified controller has been removed, or has failed. Message is generated by the partner controller.	Failure or removal of one controller (partner) in an Active-Active configuration.	Replace the controller.
Controller *x* Timeout	The specified controller has not responded in the allotted time. Message is generated by the partner controller.	Failure or removal of one controller (partner) in an Active-Active configuration.	Replace the controller.
Controller *x* Present	The specified controller has been inserted.	Partner controller has been inserted.	No action necessary.
Controller *x* Valid	The specified controller has passed its selftest and is now ready (failback).	Partner controller is ready to fail back.	No action necessary.
Controller Failover Started	The process of failing over to the remaining controller has started.	Failure or removal of the partner controller.	No action necessary.
Controller Failover Completed	The process of failing over to the remaining controller has completed.	Completion of failover process.	No action necessary.
Controller Failback Started	The process of failing back to a replacement controller has started.	Partner controller started failback.	No action necessary.
Controller Failback Completed	The process of failing back to a replacement controller has completed.	Completion of failback process.	No action necessary.
Controller Powered On	The controller has been powered on or reset.	The controller was powered on.	No action necessary.
Controller Powered Off	The controller has been powered off or removed.	Removal of controller or power.	No action necessary.
Controller Reset	The controller has been reset either locally or remotely.	User initiated a controller reset.	No action necessary.

Controller Event Messages	Description	Cause	Action
Controller Shutdown	The controller has been shutdown either locally or remotely.	User initiated a controller shutdown. The controller temperature was exceeded and the controller shut itself down.	No action necessary. Check for a failed fan, replace as needed. Check for blocked air flow, correct as needed. Check for high ambient temperature, reduce the environments ambient temperature.
Controller Selftest Passed	The controller has completed its power on selftest.	Selftest completion on startup.	No action necessary.
Controller Selftest Failed	The controller has failed its power on selftest.	Selftest failure on startup.	Replace the controller.
Controller Firmware Upgraded to version *xxxx*	The controller's firmware has been upgraded to the indicated version.	User upgraded the controller firmware.	No action necessary.
Configuration Changed	The configuration has been changed.	A change in the configuration has occurred.	If you are using the Save Configuration feature, re-save your configuration information - it no longer matches. Otherwise no action is necessary.
DMA Error	This indicates an internal error in the controller.	Internal hardware failure.	Replace the controller. Contact Technical Support.
Event Log Cleared	The controller's event log has been erased.	You have cleared the event log.	No action necessary.
Fatal Watchdog Error	The controller watchdog timer has detected a fatal hardware or firmware error.	Internal hardware or firmware failure.	Replace the controller. Contact Technical Support.

Controller Event Messages	Description	Cause	Action
Fatal SDRAM ECC	The controller has detected unrecoverable ECC errors on the SDRAM, either a multiple-bit error or uncorrectable single-bit error.	Fault SDRAM or damaged internal bus.	Replace the controller. Contact Technical Support.
Fatal Coprocessor Error	The coprocessor has failed.	Internal hardware or firmware failure on the coprocessor.	Replace the controller. Contact Technical Support.
Fatal Host Port *x* Error	The controller has detected a fatal error on the indicated host channel.	Memory or bus error on the indicated channel.	Replace the controller. Contact Technical Support.
Fatal Drive Port *x* Error	The controller has detected a fatal error on the indicated drive channel.	Memory or bus error on the indicated channel.	Replace the controller. Contact Technical Support.
Flush Mirrored Cache Started, *xxx* Entries, *yyy* MB	After failover, mirrored writeback cache needs to be flushed to the drives. This indicates the start of that process.	Failure or removal of the partner controller.	No action necessary.
Flush Mirrored Cache Completed	The flushing of mirrored writeback cache after a controller failure has completed.	Completion of mirrored cache flushing.	No action necessary.
Flush Cache Started, *xxx* Entries, *yyy* MB	If the controller is powered off with writeback cache still present, it is necessary to flush this data when power is restored. This indicates the start of that process.	Failure of power with writeback cache present.	No action necessary.
Flush Cache Completed	The flushing of writeback cache after a power failure has completed.	Completion of cache flushing.	No action necessary.
Recovered SDRAM ECC Error, Address: *xxxxxxxx* Bit: *yy*	The controller has detected and corrected a single bit ECC error in the SDRAM. The address and error bit are indicated.	SDRAM error.	If it repeats, replace the controller.

Controller Event Messages	Description	Cause	Action
Resync Started	A stripe synchronization of a RAID 5/50 set has started. This will be done when a controller fails or when powered off during a write.	A controller fails or is powered off during a RAID 5/50 write operation.	No action necessary.
Resync Completed	A stripe synchronization of a RAID 5/50 set has completed.	A controller fails or is powered off during a RAID 5/50 write operation.	No action necessary.
UPS Failure	This indicates that an attached UPS (if configured) has failed.	UPS has been disconnected or communications with the UPS has failed.	Check the UPS monitoring connections. Check the UPS, if necessary replace it.
UPS OK	This indicates that an attached UPS (if configured) is functioning correctly.	UPS is connected.	No action necessary.
Voltage Error: *Monitored Voltage: xx.yy*V	This indicates that the specified voltage has exceeded limits. This is a serious error, and should be reported to Technical Support.	Voltage regulator hardware failure. Enclosure 5V or 12V problem in the power supply.	Replace the controller. Contact Technical Support. Replace the defective power supply.

Drive events

These events are related to the drive bus and arrays.

Drive Event Messages	Description	Cause	Action
Array xx Critical	This indicates the specified array is critical.	Drive removal or failure.	Replace the hard drive and rebuild the array.
Array *xx* Expansion Started	Expansion of the indicated array has started.	Expansion has started.	No action necessary.
Array *xx* Expansion Restarted	Expansion of the indicated array has restarted, after a power cycle or failover/failback.	Expansion has restarted.	No action necessary.
Array *xx* Expansion Complete	Expansion of the indicated array has completed.	Expansion has completed.	No action necessary.
Array *xx* Initialization Started	Initialization of the indicated array has started.	Initialization has started.	No action necessary.
Array *xx* Initialization Complete	Initialization of the indicated array has completed.	Initialization has completed.	No action necessary.
Array *xx* Parity Check/Rewrite/Check and Rewrite Started	These events indicate that a RAID 5/50 Parity Check, Rewrite or Check and Rewrite has started.	Parity check started.	No action necessary.
Array *xx* Parity Check/Rewrite/Check and Rewrite Completed, yy Errors	These events indicate that a RAID 5/50 Parity Check, Rewrite, or Check and Rewrite has completed. Any errors are indicated.	Parity check completed.	No action necessary.
Array *xx* Parity Check/Rewrite/Check and Rewrite Aborted, *yy* Errors	These events indicate that a RAID 5/50 Parity Check, Rewrite, or Check and Rewrite has been stopped. Any errors up to this point are indicated.	Parity check canceled by the user.	No action necessary.

Drive Event Messages	Description	Cause	Action
Data Underrun WWN:*xx xx xx xx xx xx* P*y* OP:*zz*	The controller has detected a data underrun from the indicated drive. This is caused by the controller detecting a bad CRC in a frame, and usually indicates a link problem, either with cabling or an enclosure. The meaning of each field is as follows: WWN: Drive World Wide Name P: Drive port number OP: SCSI command code	Signal error.	Check cabling, Disk I/O card, and make sure that the hard drive is correctly seated.
Drive Busy SN:x*x xx xx xx xx xx* P*y* OP:*zz*	The drive indicated has returned a Busy status to a command. The meaning of each field is as follows: SN: Serial Number P: Drive port number OP: SCSI command code	Drive is busy and cannot accept any more commands.	No action necessary.
Drive CC SN:*xx xx xx xx xx xx* P*y* OP:*zz* SN:*aa* EXT:*bb* QL:*cc*	The drive indicated has returned a Check Condition status to a command. The meaning of each field is as follows: SN: Serial Number P: Drive port number OP: SCSI command code SN: SCSI sense key EXT: SCSI extended sense key QL: SCSI extended sense key qualifier	Typically because of a non-recoverable media error, hardware error.	No action necessary.

Drive Event Messages	Description	Cause	Action
Drive CC Failure SN:*xx xx xx xx xx xx* A:*yy* D:*zz*	The drive listed has failed because of an unrecoverable error. The meaning of each field is as follows: SN: Serial Number A: Array number D: Drive number in array	Typically because of a non-recoverable media error or hardware error.	Replace the hard drive.
Drive Inserted SN:*xx xx xx xx xx xx xx* ID*;yyy*	The drive indicated has been inserted. The meaning of each field is as follows: SN: Drive World Wide Name ID: Drive hard ID	Drive was inserted.	No action necessary.
Drive Invalid Block Size (Size) SN:*xx xx xx xx xx xx* P*y*	This indicates that the specified drive has an invalid block size, and so cannot be used. A low level format is required. The meaning of each field is as follows: SN: Serial Number P: Drive port number Size: Current block size (in hexadecimal). The controller requires 512 Bytes.	The drive has an invalid block size.	Replace the hard drive.
Drive Missing SN:*xx xx xx xx xx xx* A:*yy* D:*z*	The drive listed has been marked as failed because of being removed from the hard drive enclosure. The meaning of each field is as follows: SN: Serial Number A: Array number D: Drive number in array	Drive has been removed or bypassed by the user, or has a serious hardware error. Removal of cables connecting the enclosures. Removal of power to daisy-chained enclosure(s).	Replace the hard drive. Replace the cables. Restore power to the daisy-chained enclosure(s).

Drive Event Messages	Description	Cause	Action
Drive Rebuild Failure SN:*xx xx xx xx xx xx* A:*yy* D:*zz*	Rebuild has failed because of an unrecoverable error on another drive in the array. The meaning of each field is as follows: SN: Serial Number A: Array number D: Drive number in array	Typically because of a non-recoverable media error or hardware error.	Backup all data and restore to a new array.
Drive Status aa SN:*xx xx xx xx xx xx* P*y* OP:*zz*	The drive indicated has returned an unknown status to a command. The meaning of each field is as follows: SN: Serial Number P: Drive port number OP: SCSI command code	Unknown status returned by the hard drive.	Contact Technical Support and provide them with a copy of the event log.
Drive Task Full SN:*xx xx xx xx xx xx* P*y* OP:*zz*	The drive indicated has returned a Task Full to a command. The meaning of each field is as follows: SN: Serial Number P: Drive port number OP: SCSI command code	Drive is busy and cannot accept any more commands.	No action necessary.
Drive Timeout SN:*xx xx xx xx xx xx* P*y* OP:*zz*	The drive indicated has timed out for the specified command. The meaning of each field is as follows: SN: Serial Number P: Drive port number OP: SCSI command code	Drive hardware error.	Check cabling, Disk I/O cards, and make sure the hard drives are correctly seated.
Drive Timeout Failure SN:*xx xx xx xx xx xx* A:*yy* D:*zz*	The drive listed has failed because of a timeout. The meaning of each field is as follows: SN: Serial Number A: Array number D: Drive number in array	Drive error.	Replace the hard drive.

Drive Event Messages	Description	Cause	Action
New Drive Rebuild Failure SN:*xx xx xx xx xx xx* A:*yy* D:*zz*	Rebuild has failed because of an unrecoverable error on the new drive. The meaning of each field is as follows: SN: Serial Number A: Array number D: Drive number in array	Typically because of a non-recoverable media error, or hardware error.	Replace new drive and initiate a rebuild.
Rebuild Aborted A:*xx* D:*yy*	A rebuild has been stopped by the user on the indicated drive. The meaning of each field is as follows: A: Array number D: Drive number in array	A rebuild was canceled by the user.	No action necessary.
Rebuild Started SN:*xx xx xx xx xx xx* A:*yy* D:*zz*	A rebuild has started on the indicated drive. The meaning of each field is as follows: SN: Serial Number A: Array number D: Drive number in array	A rebuild has started.	No action necessary.
Rebuild Restarted SN:*xx xx xx xx xx xx* A:*yy* D:*zz*	A rebuild has restarted on the indicated drive. The meaning of each field is as follows: SN: Serial Number A: Array number D: Drive number in array	A rebuild has started.	No action necessary.
Rebuild Complete A:*xx* D:*yy*	A rebuild has completed on the indicated drive. The meaning of each field is as follows: A: Array number D: Drive number in array	A rebuild has completed.	No action necessary.

Drive Event Messages	Description	Cause	Action
SN:*xx xx xx xx xx xx* ID:aaa FW Download Start Rev:*yyyy*	The drive listed has started a firmware upgrade to the indicated revision. The meaning of each field is as follows: SN: Serial Number ID: Drive Target ID Rev: New firmware version	A firmware upgrade has started.	No action necessary.
SN:*xx xx xx xx xx xx* ID:*aaa* FW Download Complete	The drive listed has completed a firmware upgrade. The meaning of each field is as follows: SN: Serial Number ID: Drive Target ID	The firmware upgrade is complete.	No action necessary.
SN:*xx xx xx xx xx xx* ID:*aaa* Reallocate LBA *yyyyyyyy*	The controller has reallocated blocks at the indicated LBA in the drive. This is due to a verify failure. The meaning of each field is as follows: SN: Serial Number ID: Drive Target ID LBA: Address of reallocated blocks (hexadecimal)	Disk scrubbing detected an error.	No action necessary.
SATA Device Error: Port <yyy>	SATA link error.	Drive or SATA link error.	No action necessary.
A drive w/SN:<xx xx xx xx xx xx> has been removed	Drive Removed.	Drive has been physically removed.	No action necessary.

Controller drive port (host) events

These events are related to the host side Controller Port.

Controller Port Loop Event Messages	Description	Cause	Action
C*x*P*x* Detected Power-on/Reset at ID:*yyy* LUN:*zzz*	The specified host has accessed a logical drive for the first time, or for the first time following a reset. The meaning of each field is as follows: CxPx: Controller and host port number ID: Port ID of the host LUN: LUN number accessed	First access by a particular host after a reset.	No action necessary.
Host P*x* WWN: *yy yy yy yy yy yy* Logged in at ID:*zzz*	The indicated host system has logged into the controller. These events will only be listed for hosts that have SAN mappings created. P: Host port number WWN: Host World Wide Name ID: Port ID of the Host	Host systems logs into the controller.	No action necessary.

Controller Port Loop Event Messages	Description	Cause	Action
CC to Host ID:*xxx* LUN:*yyy* P*z* OP:*aa* SN:*bb* EXT:*cc* QL:*dd*	This indicates that an unrecoverable drive error has occurred for a particular command. This may be due to a drive error in a non fault tolerant mode, such as RAID 0, or when the array is already in a degraded mode. The controller will pass the status from the drive back to the host system, to allow the host recovery mechanisms to be used. The meaning of each field is as follows: ID: Host ID LUN: LUN requested P: Host port number OP: SCSI command code SN: SCSI sense key EXT: SCSI extended sense key QL: SCSI extended sense key qualifier	Data is not recoverable.	Backup all data, and restore to a new array.

Enclosure events

These events are related to the SES monitored enclosure components.

SES Event Messages	Description	Cause	Action
Encl.*yy* Power supply *zz* OK WWN: *xx xx xx xx xx xx xx xx*	The controller has detected that the indicated power supply is functioning correctly.	Normal condition reported.	No action necessary.
Encl.*yy* Power supply *zz* Critical WWN: *xx xx xx xx xx xx xx xx*	The controller has detected that the indicated power supply is not functioning.	The specific power supply has failed. The specific power supply is powered off.	Replace the power supply. Make sure that the specific power supply On/Off button is in the On position (|).
Encl.*yy* Power supply *zz* Not Present WWN: *xx xx xx xx xx xx xx xx*	The controller has detected that the indicated power supply is not present.	The power supply was removed.	Re-insert the power supply, connect the power cord, and power on the power supply.
Encl.*yy* Fan *zz* OK WWN: *xx xx xx xx xx xx xx xx*	The controller has detected that the indicated fan is functioning correctly.	Normal condition reported.	No action necessary.
Encl.*yy* Fan *zz* Critical WWN: *xx xx xx xx xx xx xx xx*	The controller has detected that the indicated fan is not functioning.	A specific fan failure. Total fan failure. Cooling fan module was removed.	Replace the cooling fan module. Replace the cooling fan module. Re-inset the cooling fan module.
Encl.*yy* Temp. *zz* <*aa*>C OK WWN: *xx xx xx xx xx xx xx xx*	The controller has detected that the indicated temperature sensor is within limits.	Temperature sensors are reporting normal temperatures in the enclosure.	No action required.

SES Event Messages	Description	Cause	Action
Encl.*yy* Temp. *zz* <*aa*>C Warning WWN: *xx xx xx xx xx xx xx xx*	The controller has detected that the indicated temperature sensor has reached a warning limit.	Temperature sensors are reporting enclosure temperatures have reached the threshold of 70°C.	Automatic system shutdown will begin.
			Make sure that both cooling fans are operating normally. (Replace if needed.) If the fans are set to automatic speed control, place the jumper on the Cooling fan module circuit board to force the fans to high speed. If the environment ambient temperature is high, reduce the ambient temperature. Make sure that the airflow is not blocked or restricted on the enclosure.

SES Event Messages	Description	Cause	Action
Encl.*yy* Temp. *zz* <*aa*>C Critical WWN: *xx xx xx xx xx xx xx xx*	The controller has detected that the indicated temperature sensor has reached a critical limit.	Temperature sensors are reporting enclosure temperatures have reached the threshold of 50°C.	If the fans are set to automatic speed control, place the jumper on the cooling fan module circuit board to force the fans to high speed. If the environment ambient temperature is high, reduce the ambient temperature. Make sure that the airflow is not blocked or restricted on the enclosure.
Encl.*xx* Alarm *zz* is OFF WWN: *yy yy yy yy yy yy*	All conditions are normal or the alarm is off.	No condition being reported. Alarm silenced.	No action necessary. User pressed the **Alarm Silence** button on the front panel.
Encl.*xx* Alarm *zz* is ON WWN: *yy yy yy yy yy yy*	Alarm is on.	A condition caused the alarm to sound.	Press the **Alarm Silence** button on the front panel and isolate the cause of the alarm.

Failed drives

The controller maintains a list of failed drives. Drives are listed in the following format:

Failed Drive:xx WWN: yy yy yy yy yy yy

Reason Code

The reason code can be one of the following:

Reason Codes	Reason	Action
Drive Timeout	The drive has either timed out or been removed.	Re-insert the hard drive. Replace the hard drive.
Command: xx Sense Key: *yy* Ext Sense: *zz*	The drive has failed for the specified command, with the indicated SCSI sense key and extended sense key.	Replace the hard drive.

Clearing event logs

To clear event logs:

1 From the Main Menu, select **Event Logs Menu**, then press ENTER. The *Select Controller Event Log Menu* opens.

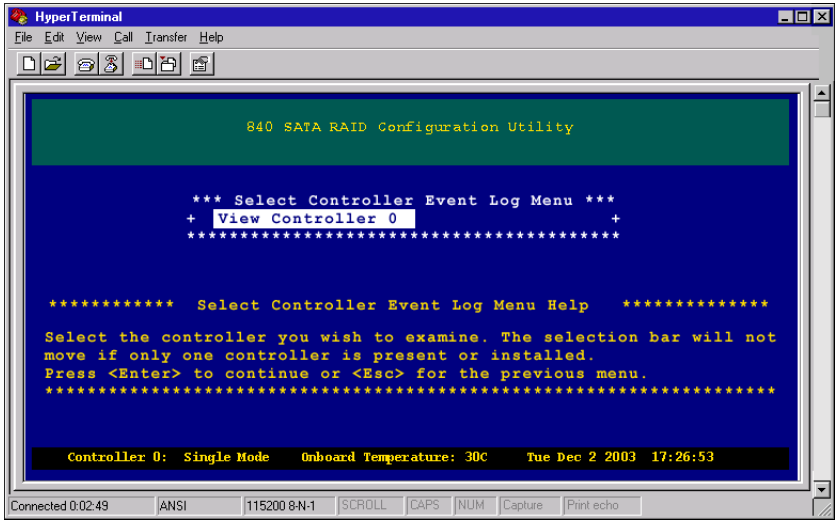

2 From the Select Controller Event Log Menu, select the controller for which you want to clear the event logs, then press ENTER. The *Event Log Types* menu opens.

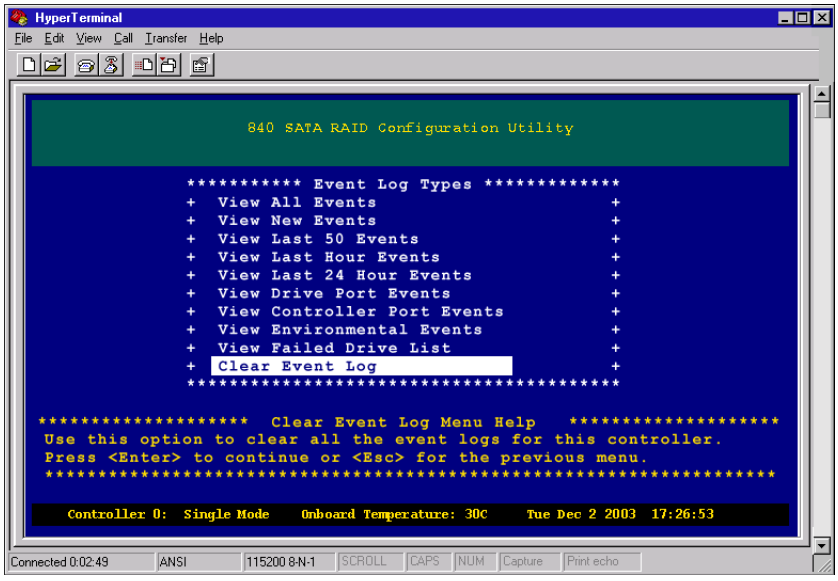

3 From the Event Log Types menu, select **Clear Event Log**, then press ENTER. The *Clear Event Log* screen opens.

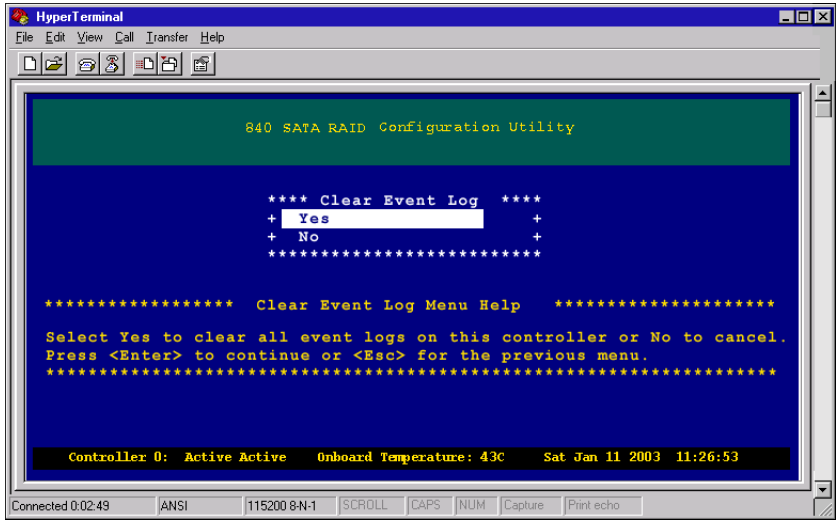

4 From the Clear Event Log screen, select **Yes** to clear all events on this controller, then press ENTER. The event log is cleared.

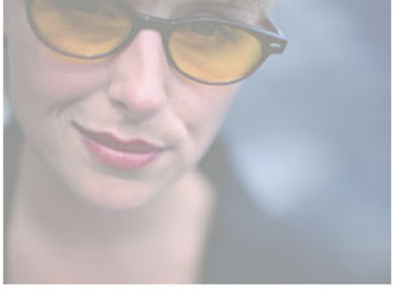

10 Diagnostics

This chapter provides information on diagnostics. Read this chapter to learn how to:

- Access diagnostics
- Use offline diagnostics
- Use online diagnostics
- Controller maintenance - downloading firmware
- Drive maintenance
- Shutdown the controller
- Use a diagnostics dump

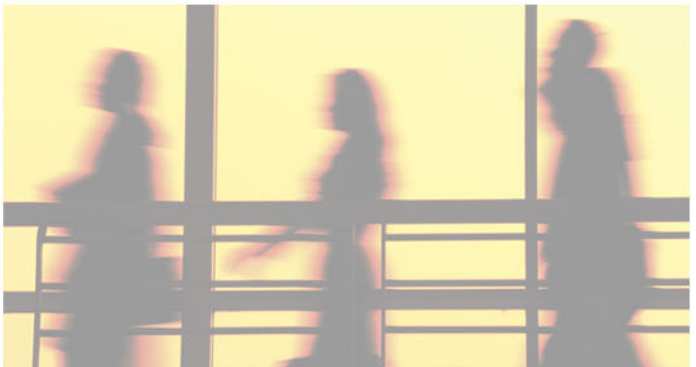

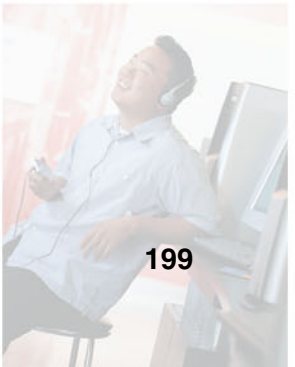

Accessing diagnostics

The RAID controller has a comprehensive VT-100 Diagnostics menu. The following operations can be performed from the diagnostics menu: Download and Program Operational Firmware, Download and Program Drive Firmware, RAID 5/50 Parity Check, Command Tracing, Shutting down the controllers, and capturing the diagnostics data.

To access diagnostics:

1. From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics Menu* opens.

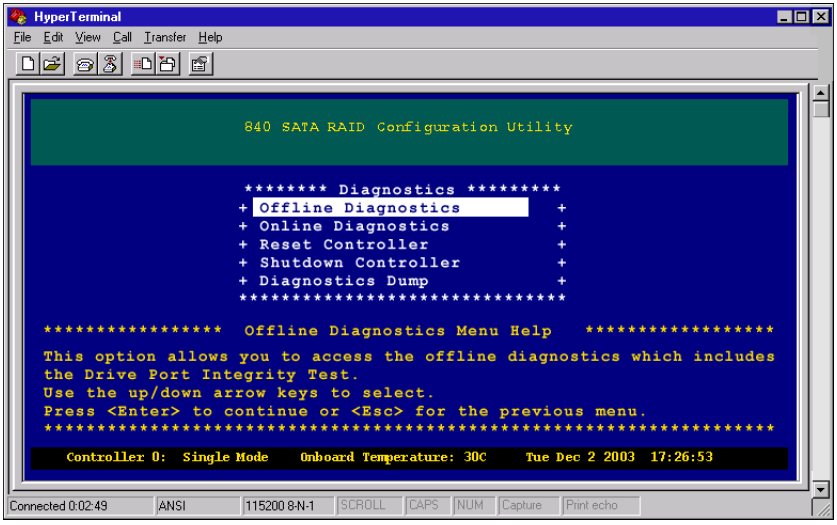

   From this menu you can enter offline or online diagnostics, or reset or shutdown the controller(s).

2. From the Diagnostics menu, select the type of diagnostics or controller operation you want to access, then press ENTER.

Important

Offline diagnostics can only be performed when there is no activity and requires that the enclosure be disconnected from the host system(s). See “Using offline diagnostics” on page 202

Online diagnostics can be performed while there is no activity in progress with the enclosure still connected to the host with no effect on the controller’s operating status. See “Using online diagnostics” on page 205

Using offline diagnostics

To view offline diagnostics:

1. Before performing this diagnostic, you must disconnect the host system(s) from the enclosure.
2. From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics* menu opens.

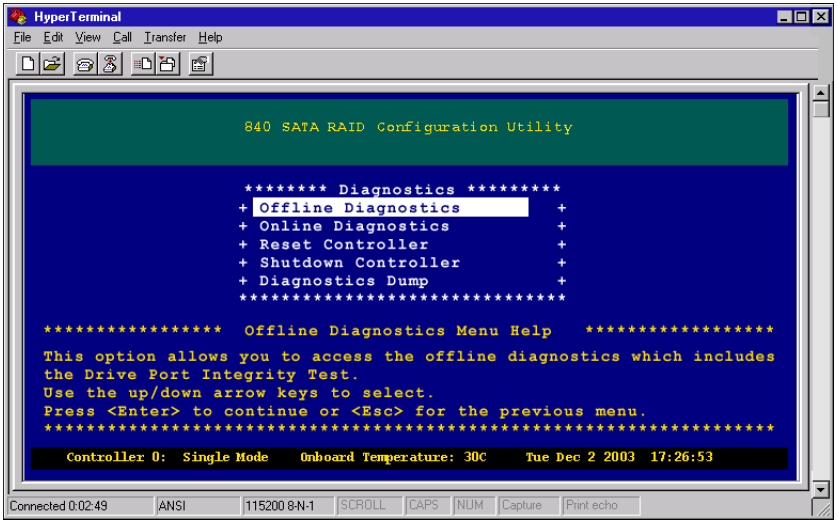

3 From the Diagnostics menu, select **Offline Diagnostics**, then press **Enter**. The *Offline Diagnostics* menu opens.

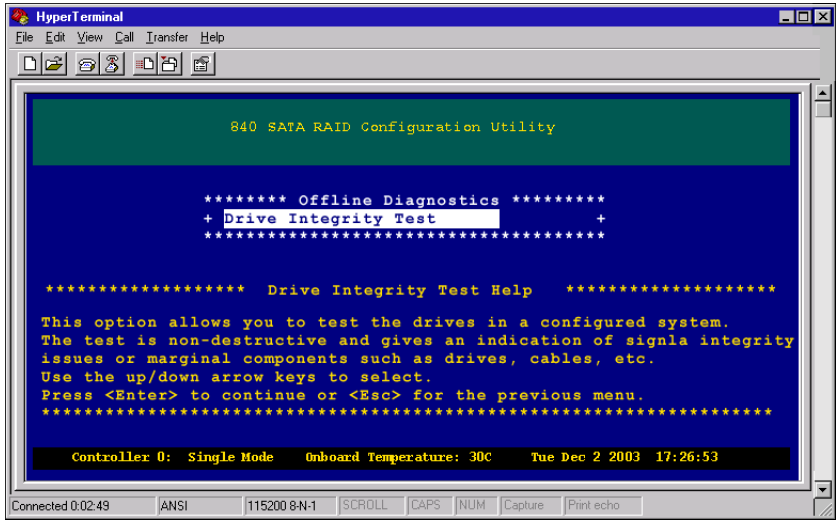

4 Press **ENTER** to continue. The *Test Data Pattern* screen opens.

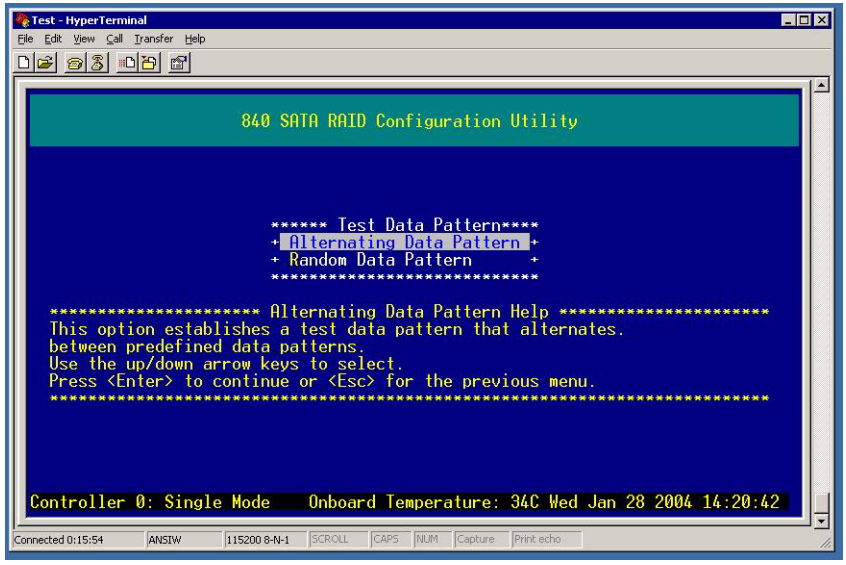

5 Select the type of test you want performed. The *Read or Write Test* screen opens.

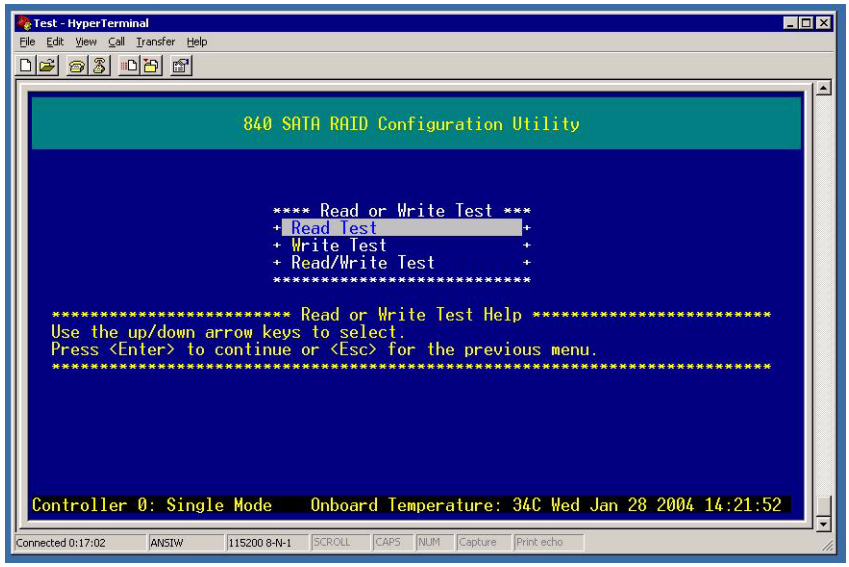

6 Select the type of test to perform, the *Test Duration* screen opens.

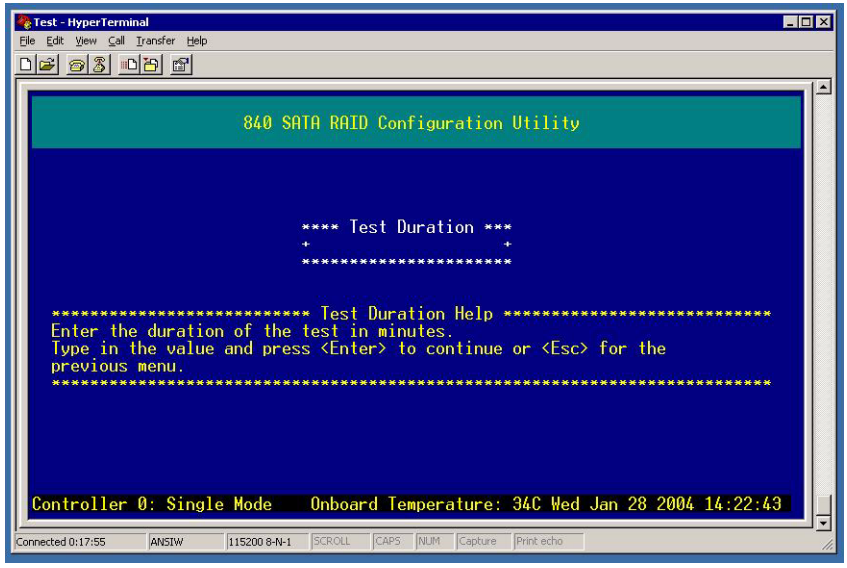

7 Type the test duration, in minutes, then press **ENTER**. The selected test is performed.

Using online diagnostics

While the controller is connected to the host, it is possible to perform Controller Maintenance, Drive Maintenance, RAID 5/50 Parity Check diagnostics, as well as to shutdown the controllers and perform a diagnostics dump to the terminal window.

To access online diagnostics:

1 From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics* menu opens.

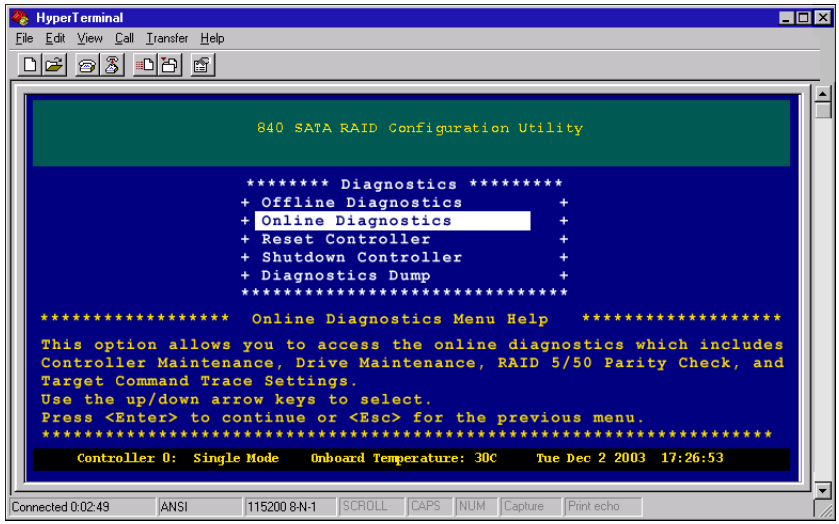

2 From the Diagnostics menu, select **Online Diagnostics**, then press ENTER. The *Online Diagnostics* menu opens.

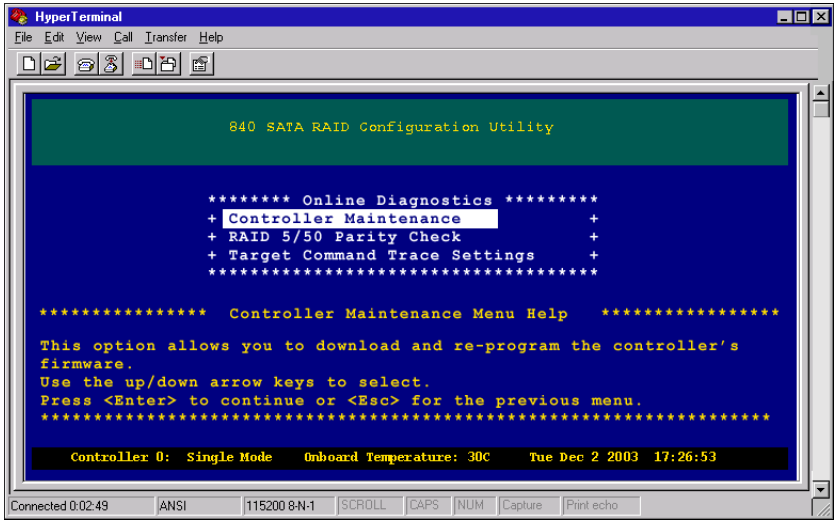

3 From the Online Diagnostics menu, select the online diagnostic test you want to perform.

Controller maintenance - downloading firmware

There may be an occasion when you need to download new or updated controller firmware or reload the current firmware. To obtain the firmware, you may need to download it from the support Web site, or get it from a CD or floppy disk.

To download and program the controller firmware:

1. Access the Online Diagnostics menu (see "To access online diagnostics:" on page 205).
2. From the Online Diagnostics menu, select **Controller Maintenance**, then press **ENTER**. The *Download and Program Controller Firmware* screen opens.

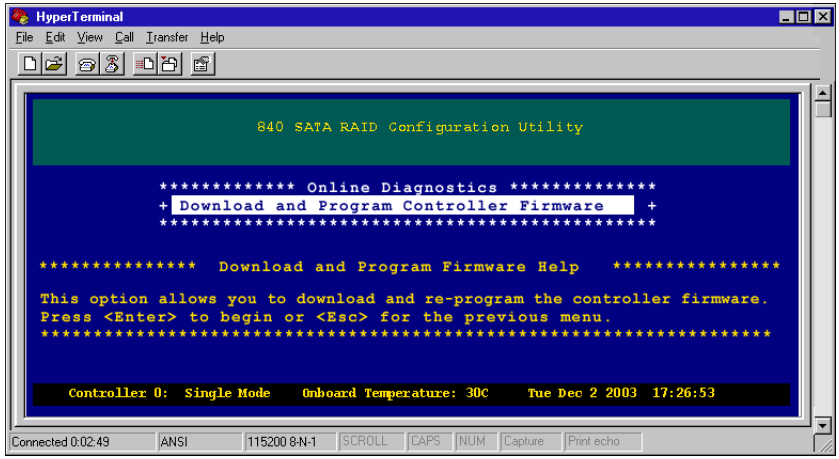

3. From the Download and Program Controller Firmware screen, press **ENTER** to begin the firmware download and re-program procedure.

   At the bottom of the screen, a series of "C" characters will be displayed.

4 When the Cs appear, select **Send File** from the **Transfer** list (not shown). The *Send File* dialog box opens.

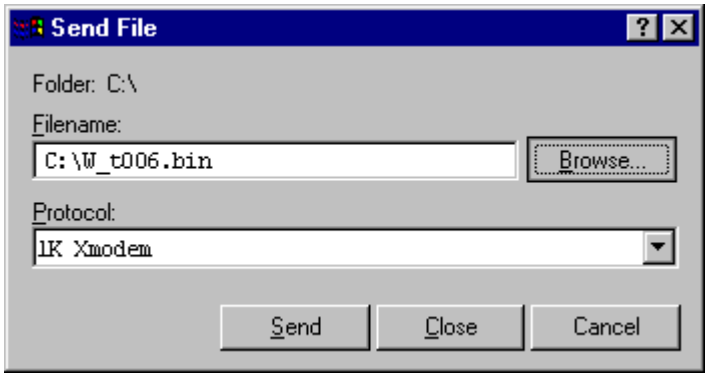

5 Click the **browse** button and locate the new firmware file (similar to "W_t006.bin"), select the **1K Xmodem** protocol, then click **Send**. The *Xmodem file send* screen opens.

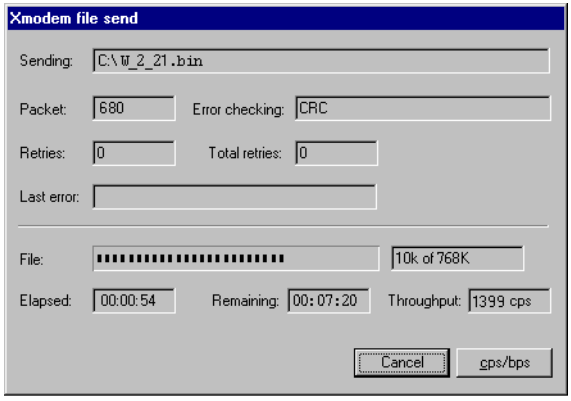

Important

Make sure that the protocol "1K Xmodem" is selected. The "1K Xmodem" protocol is preferred because of its superior error detection (CRC vs. checksum).

From the Send screen you can monitor the progress. You can safely stop the transfer without affecting your existing firmware any time during the transfer until it has been completed. The upload does not overwrite the firmware during the upload process. It writes the new code into RAM until completed, then copies the new firmware code to the EEPROM after verification.

If you elect to stop a download in progress, make sure that the stop (abort) command was completed by clicking the **Cancel** button.

6 After the transfer screen disappears, press **ENTER**. You are returned to the Main Menu.

7 Restart the controller(s).

8 From the Main Menu, select **Diagnostics** menu, then press **ENTER**. The *Diagnostics* menu opens.

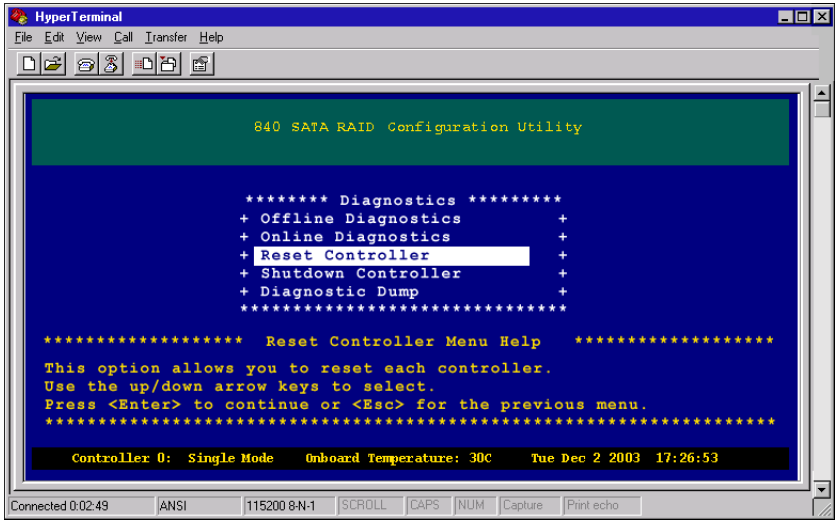

9 From the Diagnostics menu, select **Reset Controller**, then press ENTER. The *Select Controller to Reset* menu opens.

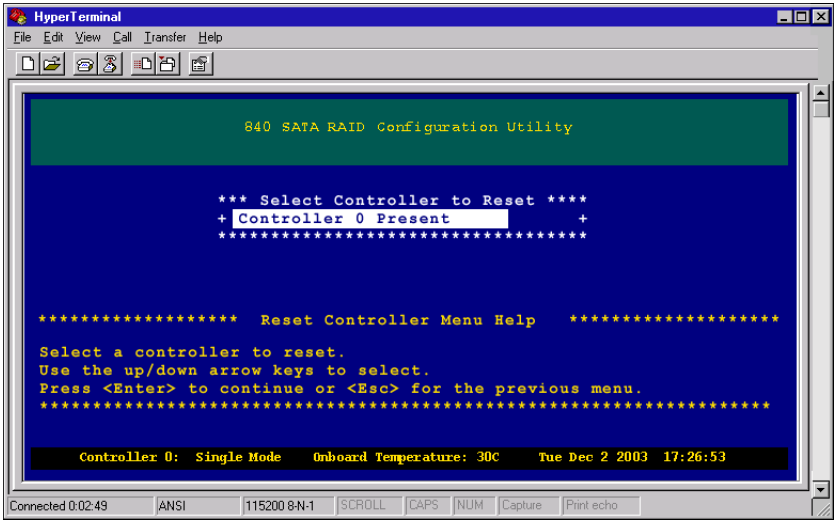

10 From the Select Controller to Reset menu, select the controller on which the firmware was changed, then press ENTER.

11 Select **Yes** to reset the controller, then press ENTER. Wait while the controller reset takes place.

The RAID Controller(s) will restart with the new firmware.

Important

In duplex mode (dual controllers), the second controller's firmware is automatically updated during the restart procedure.

Drive maintenance

Caution

We recommend that you backup data from any drives before attempting a download. During the download process, it is essential that no drives be inserted or removed, and that power is not switched off. If a drive firmware download is interrupted, it is possible that the drive will become unusable. For this reason, the controller will not download drive firmware in parallel to all drives, but one drive at a time. In this way, even if a power failure occurs, the worst case scenario is one drive will be damaged.

The controller can download firmware to the attached hard drives. Because of the large number of hard drives on the market, it is essential to use only a Gateway approved drive firmware file. Different drive manufacturers have different methods for downloading drive firmware, sometimes requiring vendor-unique utilities to perform the download. Gateway rigorously tests drive firmware and creates a single download file that can be used through the controller. This eliminates the problems associated with host-based utilities and the possibility of permanently damaging a hard drive by downloading incorrect firmware.

Download the firmware from the host

To download firmware from the host:

1 From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics* menu opens.

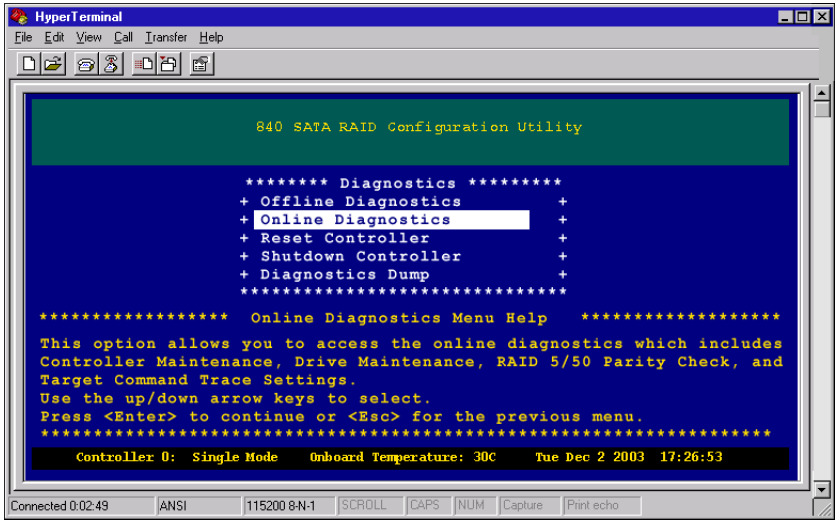

2 From the Diagnostics menu, select **Online Diagnostics**, then press ENTER. The *Online Diagnostics* menu opens.

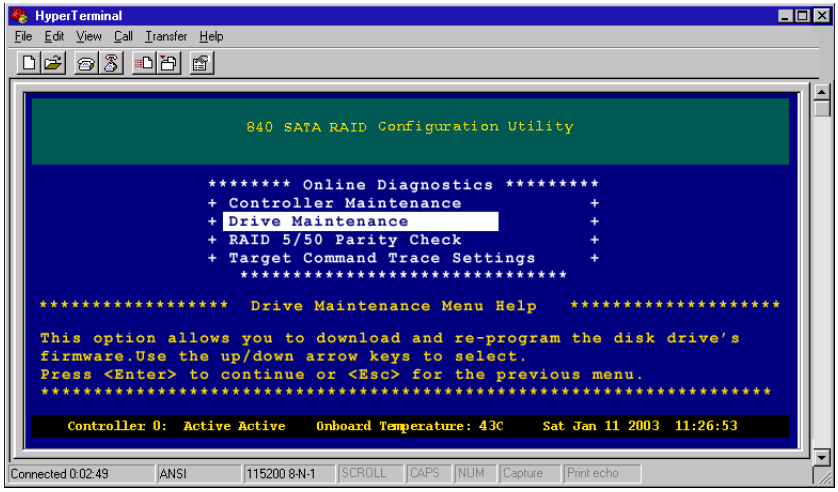

3 From the Online Diagnostics menu, select **Drive Maintenance**, then press ENTER. The *Drive Firmware Menu* opens.

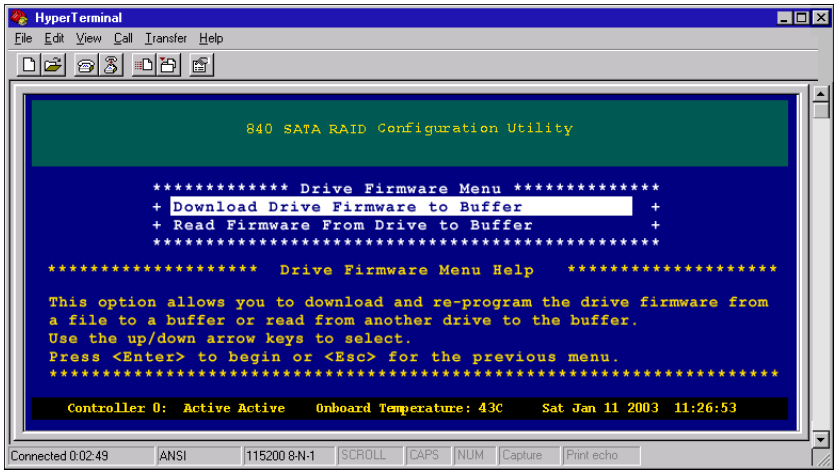

4 From the Drive Firmware Menu, select the **Download Drive Firmware to Buffer** transfer method, then press ENTER. At the bottom of the screen a series of "C" characters will be displayed.

5 When the Cs appear, select **Send File** from the **Transfer** list (not shown). The *Send File* dialog box opens.

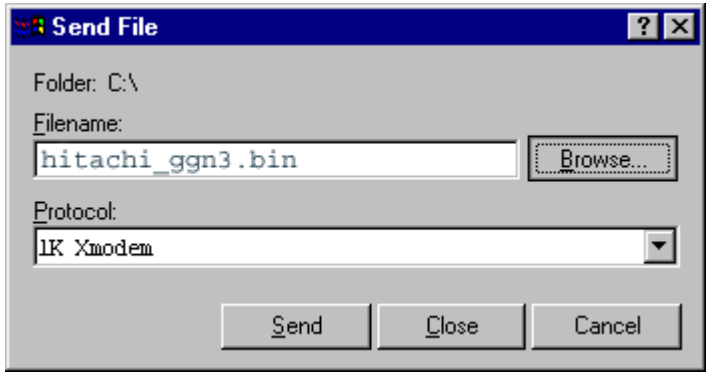

6 Click the **browse** button and locate the new firmware file (similar to "hitachi_ggn3.bin").

7 Select the **1K Xmodem** protocol, then click **Send**. The *Xmodem file send* screen opens. You can follow the progress of the download from this screen.

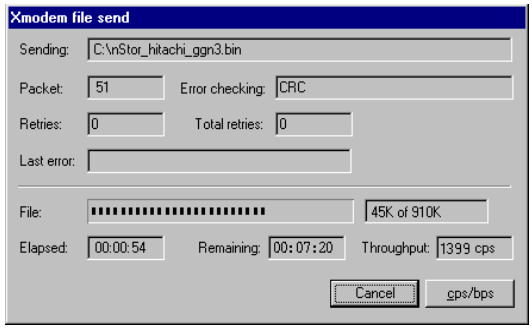

Read firmware from a drive to the buffer

To read firmware from a drive to a buffer:

1 From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics* menu opens.

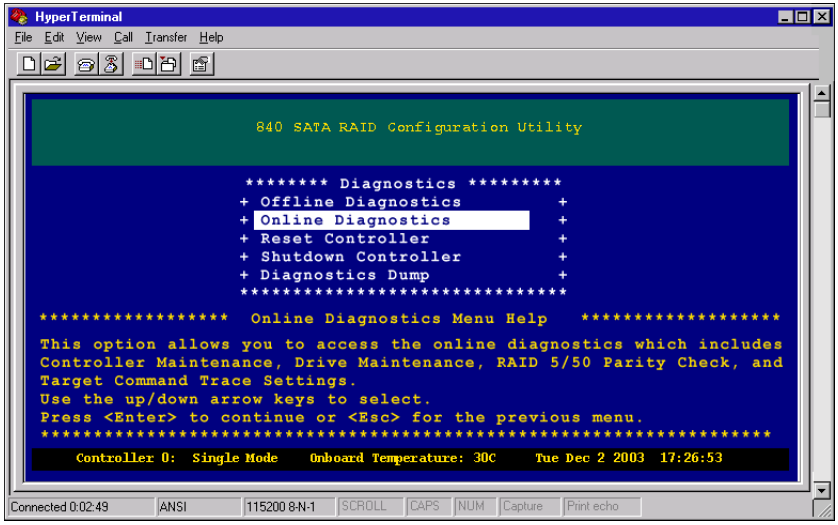

2 From the Diagnostics menu, select **Online Diagnostics**, then press ENTER. The *Online Diagnostics* menu opens.

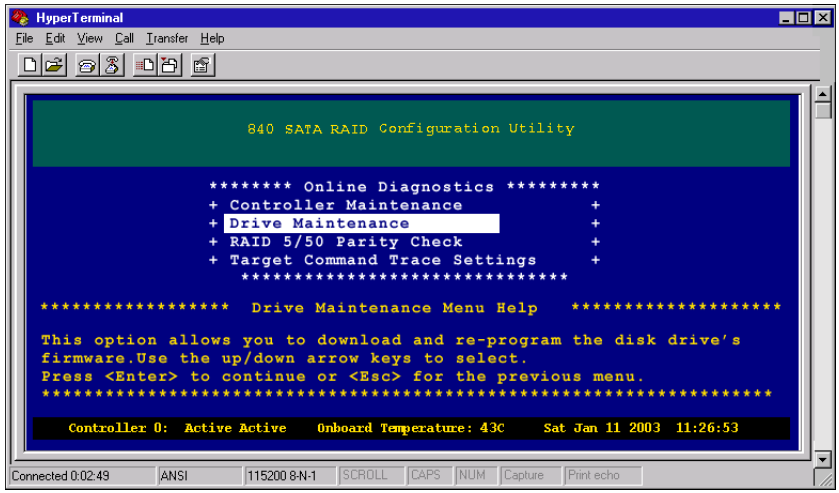

3 From the Online Diagnostics menu, select **Drive Maintenance**, then press ENTER. The *Drive Firmware Menu* opens.

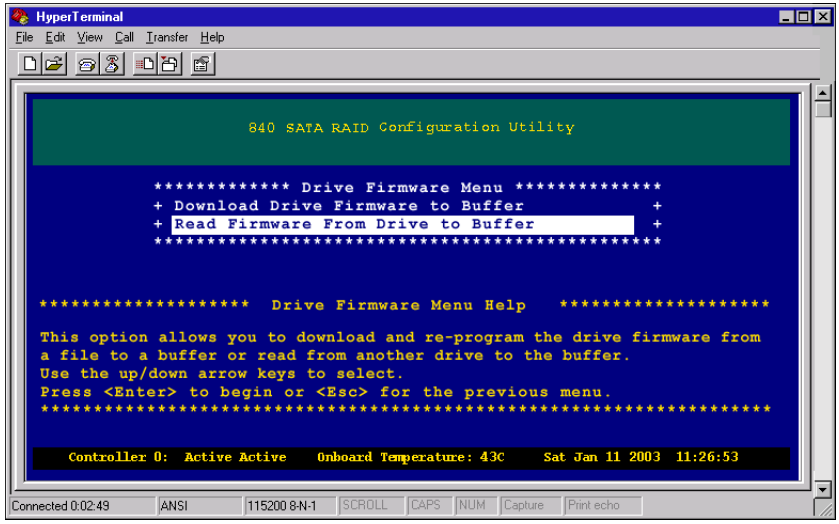

4 From the Drive Firmware Menu, select **Read Firmware From Drive to Buffer transfer method**, then press ENTER.

5 Select the appropriate firmware from the choices, then press ENTER. The *Drive Firmware* menu reopens.

6 Choose one of the available options and press ENTER.

You can download to all supported drives, all supported unused or hot spare drives, or all supported drives that are not already at this firmware level.

If any used drives are to be updated, you are warned that if a drive firmware download is interrupted, it is possible that the drive will become unusable. For this reason, the controller will not download drive firmware in parallel to all drives, but one drive at a time.

When downloading drive firmware, the controller first checks the vendor and product ID of each drive to determine whether the firmware should be updated, then it updates the drive.

7 Turn the enclosure off, then on after a drive firmware download. This will make sure that the drives use the correct firmware.

Checking RAID 5/50 parity

If an array is not initialized during creation (for example, trusted), or if both controllers are replaced after the array is turned off with write operations in progress, you may need to verify the parity data on the array. It is also prudent to perform this check periodically.

To perform the parity check:

1 From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics* menu opens.

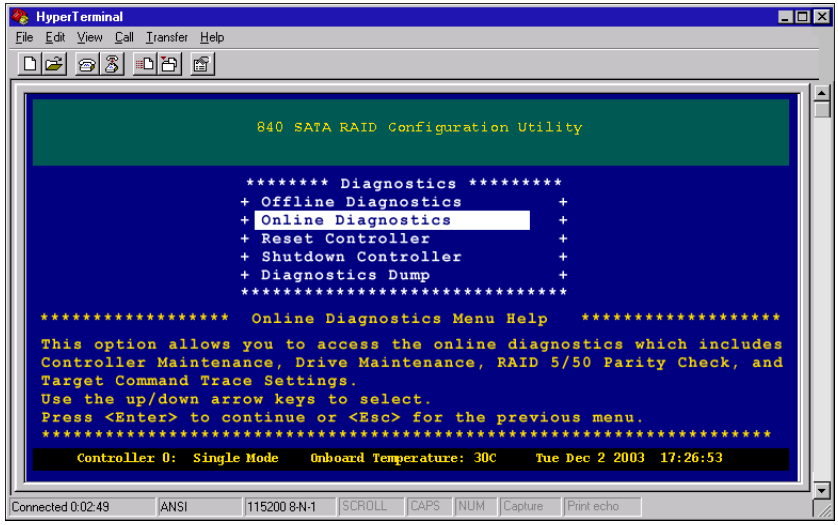

2 From the Diagnostics menu, select **Online Diagnostics**, then press **ENTER**. The *Online Diagnostics* menu opens.

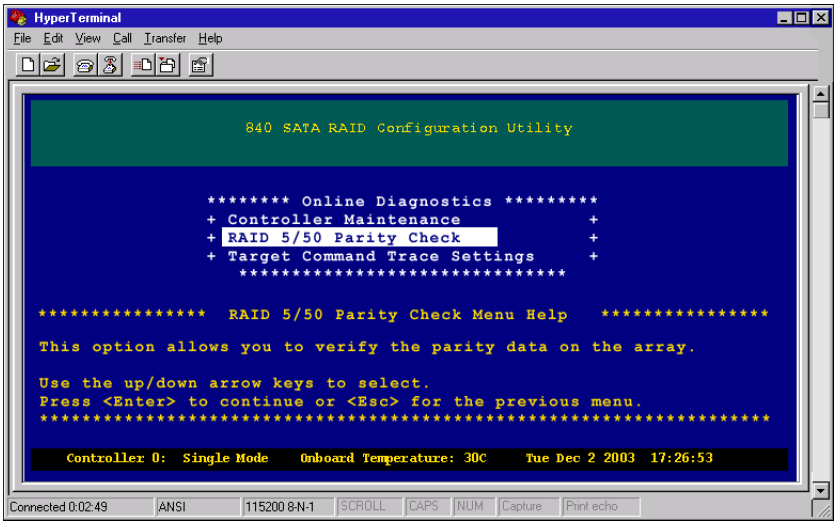

3 From the Online Diagnostics menu, select **RAID 5/50 Parity Check**, then press **ENTER**. The *Parity Check Menu* opens.

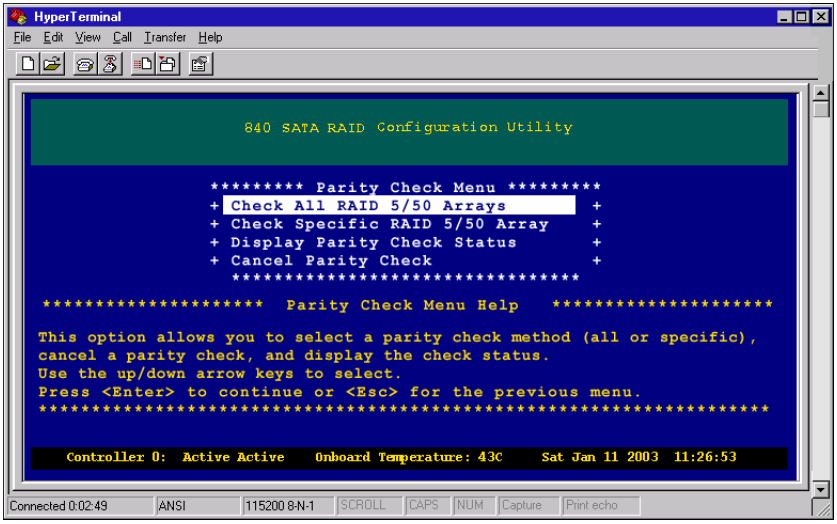

4 From the select *RAID 5/50 Parity Check Menu*, select either **All RAID 5/50 Arrays** or **Specific RAID 5/50 Arrays**, then press ENTER. The *Parity Check Option Menu* opens.

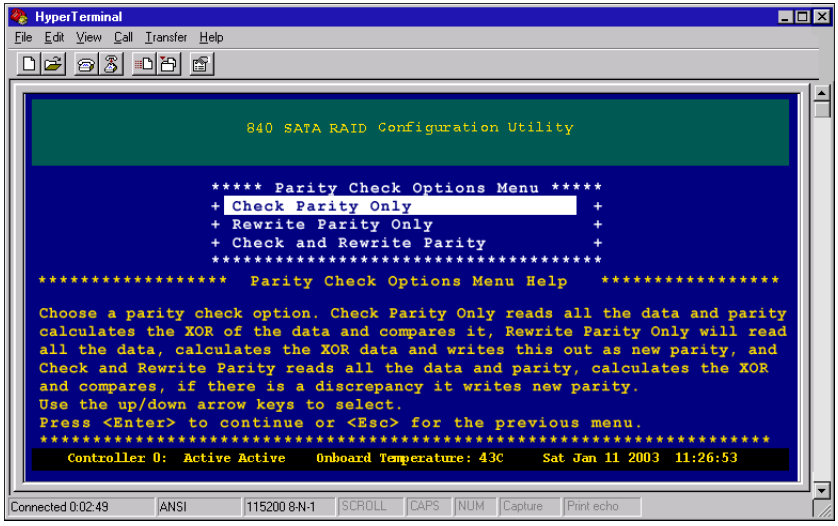

5 From the Parity Check Options Menu, select the parity check options you want to use.

Parity check options are:

Option	Description
Check Parity	This option reads all the data and parity, calculates the XOR of the data, and compares it to the parity. If there is an error, it is displayed.
Rewrite Parity	This option reads all the data, calculates the XOR of the data, and writes this out as the new parity. This is the fastest to complete, because it does not have the overhead of a comparison.
Check and Rewrite Parity	This option reads all the data and parity, calculates the XOR of the data, and compares it to the parity. Then, if there is a discrepancy, it writes out the new parity. This is the slowest to complete, because it has the overhead of a comparison as well as a rewrite.

When you have selected the required options, press **ENTER**. The *Apply Parity Check* screen opens.

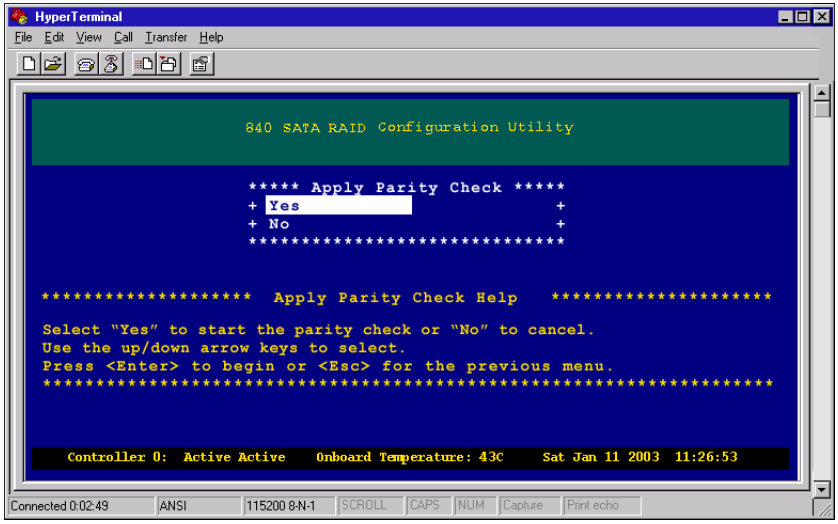

6 From the Apply Parity Check screen, select **Yes**, then press **ENTER**. The parity check starts and the *Parity Check Menu* returns.

If an error occurs during the check or rewrite, it is displayed on the terminal. Additionally, an event is entered in the event log upon start and completion.

Tracing commands

You can monitor the commands that a host is sending to the controller. This ability has a variety of uses.

It can be used to verify that a host is communicating with the controller, or as a method of debugging a system. All commands that are sent to the controller are displayed on the terminal screen.

To monitor commands:

1 From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics* menu opens.

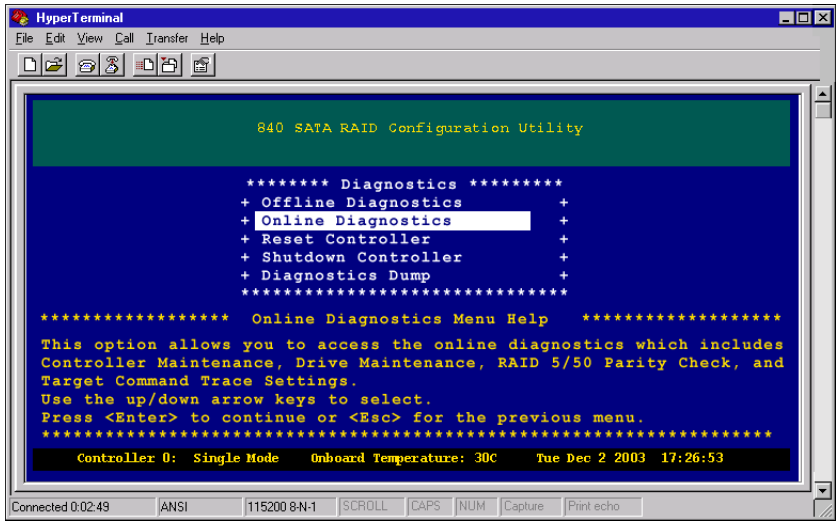

2 From the Diagnostics menu, select **Online Diagnostics**, then press ENTER. The *Online Diagnostics* menu opens.

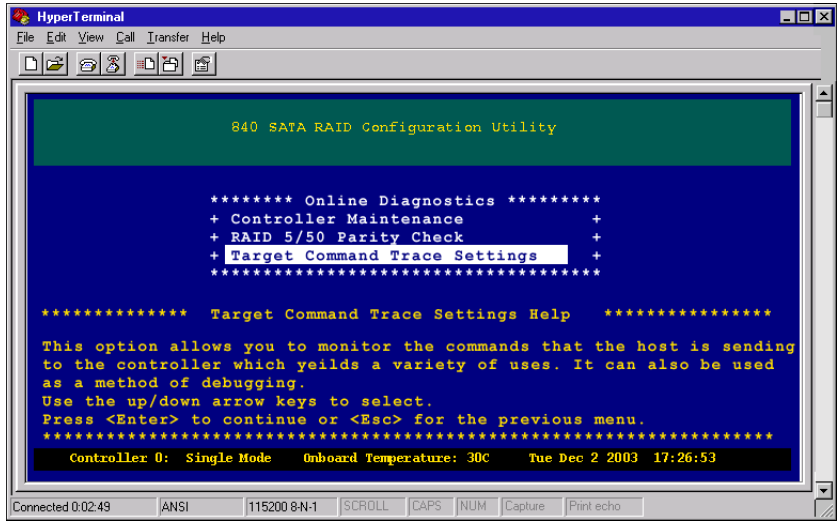

3 From the Online Diagnostics menu, select **Target Command Trace Settings**, then press **ENTER**. The *Target Command Trace Settings* screen opens.

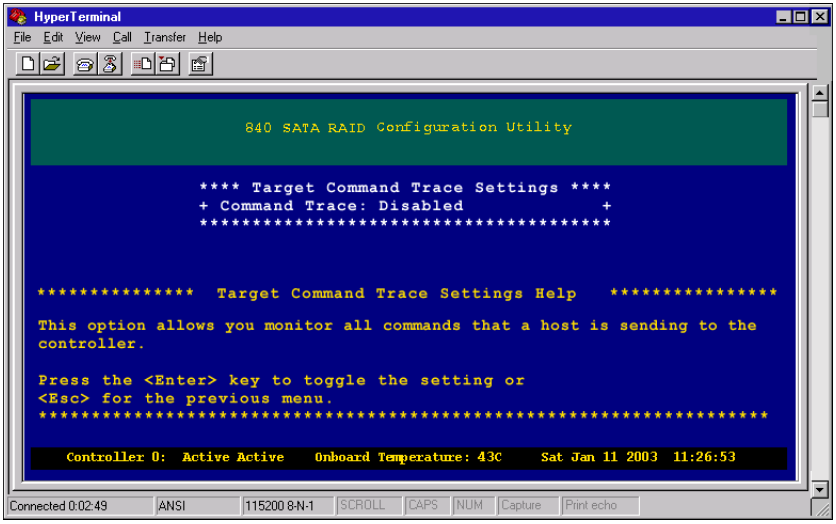

4 From the Target Command Trace Settings screen, press **ENTER** to toggle the command trace to be *Enabled* or *Disabled*.

When enabled, all incoming commands will be displayed on the terminal until the trace mode is disabled.

5 Press **ESC** to return to the previous menu.

Important

The command is on port 1 (T1) from the host with ID 124 (ID:7c) for LUN 0. The command is a SCSI inquiry command (CDB: 12).

Shutting down the controller

This option lets you gracefully shutdown the controller. When this method is used, all cache is written before the controller is shutdown.

To shutdown the controllers:

1 From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics* menu opens.

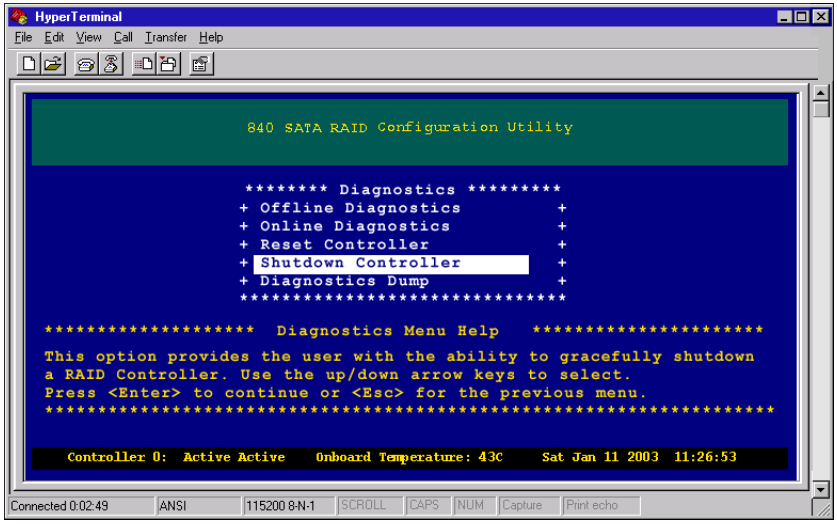

2 From the Diagnostics menu, select **Shutdown Controller**, then press ENTER. The *Select Controller to Shutdown* screen opens.

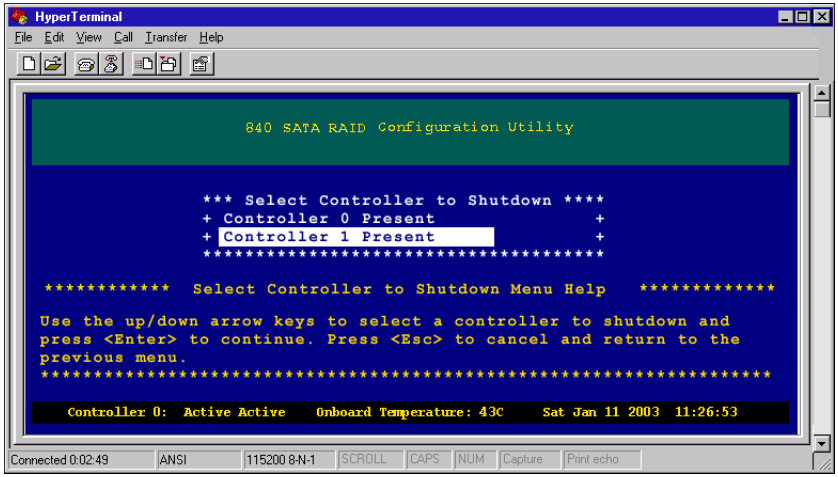

3 From the Select Controller to Shutdown screen, select a controller you want to shutdown, then press ENTER. The *Shutdown Controller* screen opens.

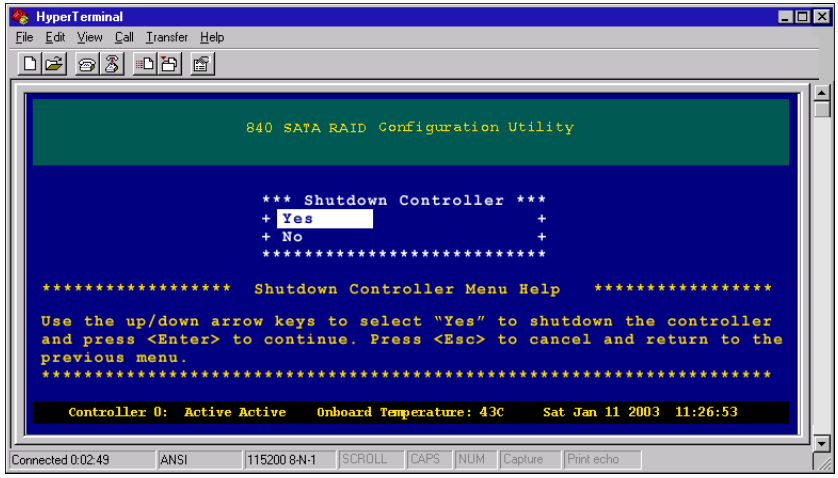

4 From the Shutdown Controller screen, select **Yes** to shutdown the controller and return to the Main Menu. The Controller will shutdown.

5 To restart the controller, you can either remove and re-insert the RAID controller in its slot or turn the enclosure off then back on.

Dumping diagnostics

During the process of troubleshooting a problem, it may be helpful to capture the diagnostics data. The information capture to the terminal screen can then be saved to an external file for later analysis.

To do a diagnostics dump:

1 From the Main Menu, select **Diagnostics Menu**, then press ENTER. The *Diagnostics* menu opens.

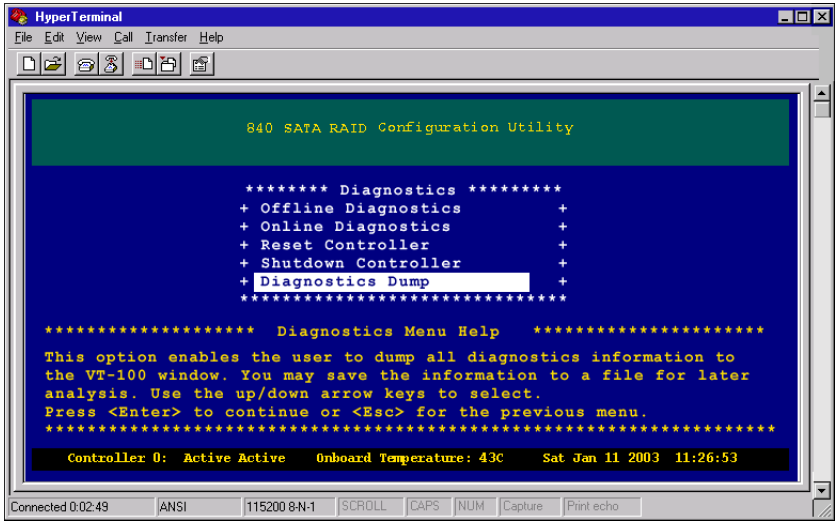

2 From the Diagnostics menu, select **Diagnostics Dump**, then press ENTER. The diagnostic data is immediately sent to the terminal window.

3 To export the data to a file, press CTRL + W to enter Text Mode.

4 Select **Capture Text** from the Transfer list, then type a file name and a destination (or note the default name and path). Click **Start** to continue.

5 Repeat steps 1 and 2.

Tips & Tricks

You will notice the cursor jumps to the lower-right-corner of the screen and after completion will return to the upper-left-corner of the screen.

6 Select **Capture Text** > **Stop** from the **Transfer** list.

7 Access the text file using a text editor.

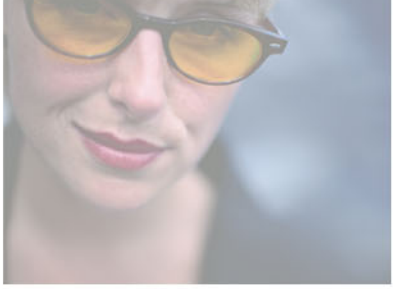

Optimizing RAID 5 Write Performance

A

This chapter provides information on optimizing RAID 5 write performance. Read this chapter to learn how to:

- Understand optimization parameters
- Use RAID 5 sub-arrays
- Promote faster rebuilds

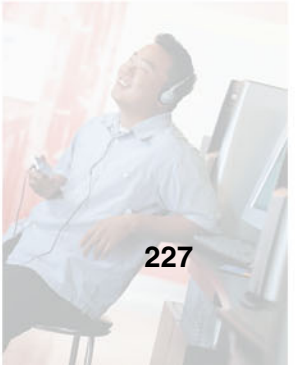

Introduction

With a typical RAID 5 implementation, there are a number of steps necessary to write the data to the media. Every write from the host system will typically generate two XOR operations and their associated data transfers to two drives. If the accesses are sequential, the parity information will be updated a number of times in succession. However, if the host writes sufficient data to cover a complete stripe, the parity data does not need to be updated for each write. It can be recalculated instead. This operation takes only one XOR operation per host write, compared to two for a standard RAID 5 write. The number of data transfers necessary are also reduced, increasing the available bandwidth. This type of write access is termed a *Full Stripe Write*.

P (20-23)	C20	C21	C22	C23	Stripe 5
C16	C17	C18	C19	P (16-19)	Stripe 4
C12	C13	C14	P (12-15)	C15	Stripe 3
C8	C9	P (8-11)	C10	C11	Stripe 2
C4	P (4-7)	C5	C6	C7	Stripe 1
P (0-3)	C0	C1	C2	C3	Stripe 0

The table shows the distribution of data chunks (denoted by Cx) and their associated parity (denoted by P(y-z)) in a RAID 5 array of five drives. An *array* is defined as a set of drives, on which data is distributed. An array will have one RAID level. A *chunk* is the amount of contiguous data stored on one drive before the controller switches over to the next drive. This parameter is adjustable from 64 K to 256 K, and should be carefully chosen to match the access sizes of the operating system. A *stripe* is a set of disk chunks in an array with the same address. In the above example, Stripe 0 consists of C0, C1, C2, and C3 and their associated parity.

Maximum performance will be achieved when all drives are performing multiple commands in parallel. To take advantage of a Full Stripe Write, the host has to send enough data to the controller. This can be accomplished in two ways. If the host sends one command with sufficient data to fill a stripe, then the controller can perform a Full Stripe Write. Alternatively, if the host sends multiple sequential commands, smaller than a stripe size (typically matching the chunk size), the controller can internally combine these commands to get the same effect. In the above example, if a 256 K chunk size

is used, then the stripe size is 1MB (4 chunks × 256 K). So, for maximum performance, the host could either send 5 × 1 MB write commands, or 20 × 256 K write commands.

The effectiveness of the controller's ability to perform a Full Stripe Write depends on a number of parameters.

Sequential access

If the commands sent from the host are not sequential, the controller will not be able to cluster them together. So, unless each individual access is sufficient to fill a stripe, a Full Stripe Write will not occur.

Number of outstanding commands

For the controller to successfully cluster commands, there has to be a number of write commands sent simultaneously. Setting the host to send up to 64 commands should prove adequate. Alternatively, enabling writeback cache will have a similar effect, as the controller can then cluster sequential commands even if the host only sends a small number of commands at a time.

Access size

With very small accesses, it is necessary to have a large number of commands to cluster together to fill up a full stripe. So, the larger the access size the better. It is best to use an access size that will fill a chunk. Of course, even if a stripe is not filled up, small sequential writes will still benefit from command clustering.

Access alignment

The alignment of a command from a host system is determined by the command's address. In an optimal system, a write of one chunk of data would reside exactly within a chunk on one disk. However, if this is not the case, this write will be split up into two separate writes to two different data drives. This will have a negative effect on performance. To overcome these problems, you can, with more sophisticated operating systems, set the access size and alignment to an optimal value.

To get the highest performance from this system, it is necessary to have a number of stripes being written in parallel. As the array expands, with more and more drives, the number of commands (and amount of sequential data) necessary to do this increases.

P (56-62)	C56	C57	C58	C59	C60	C61	C62	Stripe 8
C49	C50	C51	C52	C53	C54	C55	P (49-55)	Stripe 7
C42	C43	C44	C45	C46	C47	P (42-48)	C48	Stripe 6
C35	C36	C37	C38	C39	P (20-23)	C61	C62	Stripe 5
C28	C29	C30	C31	P (28-34)	C32	C33	C34	Stripe 4
C21	C22	C23	P (21-27)	C24	C25	C26	C27	Stripe 3
C14	C15	P (14-20)	C16	C17	C18	C19	C20	Stripe 2
C7	P (7-13)	C8	C9	C10	C11	C12	C13	Stripe 1
P (0-6)	C0	C1	C2	C3	C4	C5	C6	Stripe 0

In this table we can see that seven chunks of sequential data are necessary to fill a stripe. To have multiple commands active for all hard drives requires more data than with five drives. As can be seen, this number will increase as the number of drives increases. If a large number of drives are used, it can get difficult to achieve maximum performance, because it becomes more difficult to cluster a large number of commands to achieve a Full Stripe Write.

Using RAID 5 sub-arrays

The difficulty in realizing the maximum performance possible introduces the concept of a Sub-Array. Suppose an array consisted of two RAID 5 sets. If these are then striped, the resulting array would appear as shown in the second of the following tables. In this case, in order for a Full Stripe Write to be performed,

it is still only necessary to cluster four write commands together, as opposed to the seven necessary as indicated below. The array of drives appears as two separate sub-arrays, each with its own rotating parity.

P (40-43)	C40	C41	C42	C43	P (44-47)	C44	C45	C46	C47	Stripe 5
C32	C33	C34	C35	P (32-35)	C36	C37	C38	C39	P (36-39)	Stripe 4
C24	C25	C26	P (24-27)	C27	C28	C29	C30	P (28-31)	C31	Stripe 3
C16	C17	P (16-19)	C18	C19	C20	C21	P (20-23)	C22	C23	Stripe 2
C8	P (8-11)	C9	C10	C11	C12	P (12-15)	C13	C14	C15	Stripe 1
P (0-3)	C0	C1	C2	C3	P (4-7)	C4	C5	C6	C7	Stripe 0

It can be seen that the more sub-arrays used, the more likely it is for a Full Stripe Write to occur. We recommend that you use either four or five drives in a sub-array, for best performance. The following table shows that even with 15 drives, it is still possible to perform Full Stripe Writes, by clustering together 4 chunks of data.

P	C60	C61	C62	C63	P	C64	C65	C66	C67	P	C68	C69	C70	C71	Stripe 5
C48	C49	C50	C51	P	C52	C53	C54	C55	P	C56	C57	C58	C59	P	Stripe 4
C36	C37	C38	P	C39	C40	C41	C42	P	C43	C44	C45	C46	P	C47	Stripe 3
C24	C25	P	C26	C27	C28	C29	P	C30	C31	C32	C33	P	C34	C35	Stripe 2
C12	P	C13	C14	C15	C16	P	C17	C18	C19	C20	P	C21	C22	C23	Stripe 1
P	C0	C1	C2	C3	P	C4	C5	C6	C7	P	C8	C9	C10	C11	Stripe 0

As well as the performance advantage gained from using multiple sub-arrays, there are a number of other advantages.

Experiencing multiple drive failures

In a configuration with multiple sub-arrays, it is possible for the array to sustain multiple drive failures, provided that there is only one failure in each sub-array.

Promoting faster rebuild

A rebuild operation must read data and calculate parity from all the remaining drives in the array. If multiple sub-arrays are used, this means that it is only necessary to read the data from the remaining drives in the sub-array, not all

of the drives in the array. This increases both the rebuild speed and the speed of access to missing data, which also has to be recreated from the remaining drives.

Summary

In summary, for maximum performance using RAID 5, we recommend that you use four or five drives in a sub-array. If there are more than five drives in a sub-array, it is better to use a smaller chunk size, say 64 K or 128 K, as this will lead to more Full Stripe Writes.

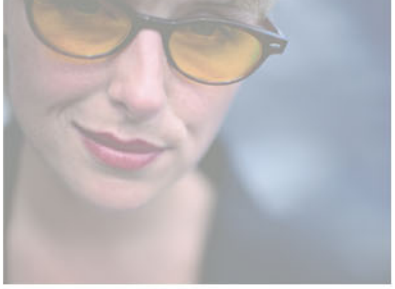

Technical Support

B

This chapter provides information on troubleshooting. Read this chapter to learn about:

- Getting Technical Support
- Notices

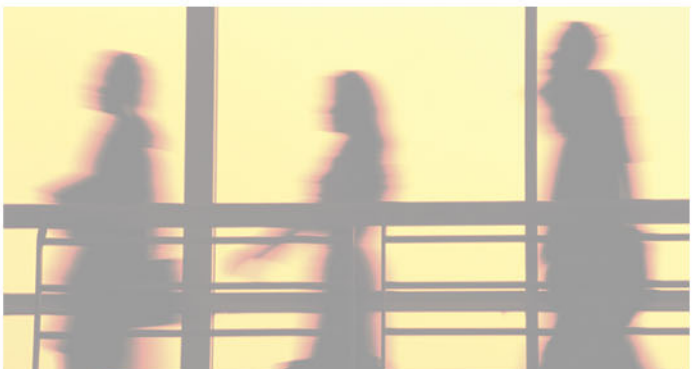

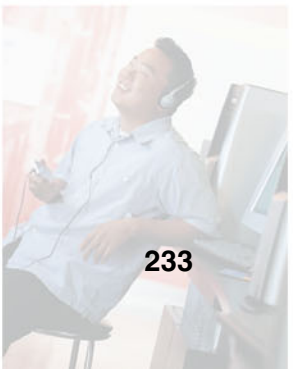

Technical Support

Gateway offers a wide range of customer service, technical support, and information services.

Telephone numbers

You can access the following services through your telephone to get answers to your questions:

Resource	Service description	How to reach
Fax on demand support	Order a catalog of documents on common problems, then order documents by document numbers. The documents will be faxed to you.	800-846-4526 (US) 877-709-2951 (Canada)
Gateway's fee-based software tutorial service	Get tutorial assistance for software issues billed by the minute.	800-229-1103 (charged to your credit card)
Gateway Technical Support	Talk to a Gateway Technical Support representative about a non-tutorial technical support question.) TDD Technical Support (for hearing impaired) is available: Weekdays 6:00 a.m. - 8:00 p.m. Central Time Weekends 6:00 a.m. - 5:00 p.m. Central Time	877-485-1464 (US) 800-846-3609 (Canada and Puerto Rico) 605-232-2191 (all other countries) 800-846-1778 (TDD)
Sales, accounting, and warranty	Get information about available systems, pricing, orders, billing statements, warranty service, or other non-technical issues.	800-846-2000 (US) 888-888-2037 (Canada)

Notices

14303 Gateway Place
Poway, CA 92064 USA

All Rights Reserved

The information in this manual has been carefully checked and is believed to be accurate. However, changes are made periodically. These changes are incorporated in newer publication editions. Gateway may improve and/or change products described in this publication at any time. Due to continuing system improvements, Gateway is not responsible for inaccurate information which may appear in this manual. For the latest product updates, consult the Gateway Web site at www.gateway.com. In no event will Gateway be liable for direct, indirect, special, exemplary, incidental, or consequential damages resulting from any defect or omission in this manual, even if advised of the possibility of such damages.

In the interest of continued product development, Gateway reserves the right to make improvements in this manual and the products it describes at any time, without notices or obligation.

Trademark Acknowledgments

Gateway and the Black-and-White Spot Design are trademarks or registered trademarks in the U.S. and other countries. SpotShop, Spotshop.com, and Your:)Ware are trademarks of Gateway, Inc. Intel, Intel Inside logo, and Pentium are registered trademarks and MMX is a trademark of Intel Corporation. Microsoft, MS, MS-DOS, and Windows are trademarks or registered trademarks of Microsoft Corporation. All other product names mentioned herein are used for identification purposes only, and may be the trademarks or registered trademarks of their respective companies.

Index

D

E

F

Q

R

S

T

U

V

W

A MAN NSTOR VT-100 SW GDE R0 1/04

8509736